Working Paper

Capacity Investment with Demand Learning

Anyan Qi
Naveen Jindal School of Management
University of Texas at Dallas

Hyun-Soo Ahn
Stephen M. Ross School of Business
University of Michigan

Amitabh Sinha
Stephen M. Ross School of Business
University of Michigan

Ross School of Business Working Paper
Working Paper No. 1321
July 2016

Capacity Investment with Demand Learning

Anyan Qi
Naveen Jindal School of Management, The University of Texas at Dallas, Richardson, TX 75080,
axq140430@utdallas.edu

Hyun-Soo Ahn, Amitabh Sinha
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI 48109,
hsahn@umich.edu, amitabh@umich.edu

We study a firm's optimal strategy to adjust its capacity using demand information. The capacity adjustment is costly and often subject to managerial hurdles which sometimes make it difficult to adjust capacity multiple times. In order to clearly analyze the impact of demand learning on the firm's decision, we study two scenarios. In the first scenario, the firm's capacity adjustment cost increases significantly with respect to the number of adjustments because of significant managerial hurdles, and resultantly the firm has a single opportunity to adjust capacity (*single adjustment scenario*). In the second scenario, the capacity adjustment costs do not change with respect to the number of adjustments because of little managerial hurdles, and therefore the firm has multiple opportunities to adjust capacity (*multiple adjustment scenario*). For both scenarios, we first formulate the problem as a stochastic dynamic program, and then characterize the firm's *optimal* policy: when to adjust and by how much. We show that the optimal decision on when and by how much to change the capacity is not monotone in the likelihood of high demand in the single adjustment scenario, while the optimal decision is monotone under mild conditions and the optimal policy is a control band policy in the multiple adjustment scenario. The sharp contrast reflects the impact of demand learning on the firm's optimal capacity decision. Since computing and implementing the optimal policy is not tractable for general problems, we develop a data-driven heuristic for each scenario. In the single adjustment scenario, we show that a two-step heuristic which explores demand for an appropriately chosen length of time and adjusts the capacity based on the observed demand is *asymptotically optimal*, and prove the *convergence rate*. In the multiple adjustment scenario, we also show that a multi-step heuristic under which the firm adjusts its capacity at a predetermined set of periods with exponentially increasing gap between two consecutive decisions is asymptotically optimal and show its convergence rate. We finally apply our heuristics to a numerical study and demonstrate the performance and robustness of the heuristics.

Subject classifications: capacity investment; demand learning; exploration-exploitation; Bayesian updating; data-driven.
Area of review: Operations and Supply Chains

1. Introduction

In most cases capacity investment requires significant time and resource commitment. Because of this, many capacity decisions are made when there exist significant demand uncertainties. While early capacity installation enables a firm to seize a time-to-market opportunity, installing capacity with little market information may result in a significant mismatch when the realized demand is seriously different from the capacity level. As choosing a "perfect" initial capacity level well before the planning horizon is often impossible, many firms adjust their capacity levels after observing some demand information during the planning horizon. For this strategy to be successful, the firm should be able to evaluate the benefit and cost between two options – waiting it out (gathering more information) and committing to an action (adjusting the capacity level) – a classic trade-off between *exploration* and *exploitation*.

However, changing capacity later on is often difficult. Changing capacity often requires a considerable amount of time. It may take several weeks or months to get new machines or workers ready for production, layoff current labor, or disinvest installed equipment. Increasing capacity by adding new machines and/or hiring new workers is expensive and often irrevocable. Downsizing the capacity level, which typically requires layoffs and equipment disinvestment, can also be costly. Moreover, the *de facto* capacity adjustment cost may increase in the number of capacity adjustments that the firm has committed to. For example, when Genentech, a biopharmaceutical company, evaluated the option to build the third cell culture production plant (CCP3) in Vacaville, CA, it was noted that the effective cost to build this plant would be higher compared to CCP2 which was being built for the managerial hurdles associated with managing this additional construction project as well as lower effective operating efficiencies caused by the diseconomy (Snow et al. 2006). As another example, in a major cosmetic company that the authors have intimate knowledge about, a product manager needs to get approval from senior executives to change capacity. As frequent requests to change capacity may be perceived as a sign of incompetence, this managerial hurdle makes it difficult to change capacity often. We model the fact that it becomes increasingly difficult to change capacity as the capacity adjustment cost increases in the number of adjustments made.

Naturally, if the capacity adjustment features a long leadtime and high capacity adjustment costs, or the capacity adjustment increases significantly with respect to the number of adjustments, less frequent capacity adjustments are observed such as the case in the biopharmaceutical industry (Kaminsky and Yuen 2014) and the major cosmetic company above. On the other hand, if the capacity adjustment features a short leadtime and low capacity adjustment costs, and the capacity adjustment does not vary significantly with respect to the number of adjustments, more frequent capacity adjustments are observed such as the case in the IT-based company, where the capacity adjustment for a data center features significantly low adjustment cost, extremely short leadtime, and no managerial hurdles. Therefore, the capacity adjustment cost is independent of the number of adjustments and a dynamic capacity adjustment strategy in real time is feasible to manage the data center capacity (*Power Assure* 2009).

Motivated by these observations, we examine a make-to-order firm's capacity decision using demand observation: ***when**, and by **how much**, should a firm adjust its capacity?* To investigate this question, we consider a firm selling a single product for a finite planning horizon when the firm has only partial information about random demand. In each period, the firm observes the realized demand and collects more information. Based on the information, the firm actively updates its knowledge about demand, and uses this updated knowledge in the capacity adjustment decision. Because of managerial hurdles in capacity adjustment, we consider the scenario where the capacity adjustment cost may increase with the number of capacity adjustments.

We first build a general model where the capacity adjustment costs increase in the number of capacity adjustment that the firm has committed to. We show that the decision problem can be equivalently characterized as one featuring a state-dependent fixed cost. This problem is in general intractable and we show that even in the special case where the fixed cost is not state-dependent, the optimal policy can be rather complicated, which clouds the impact of demand learning.

In order to clearly analyze the impact of demand learning on the firm's capacity adjustment decision, we consider two special settings of the general model, which differ in how fast the capacity adjustment cost increases with respect to the number of capacity adjustments. In the first setting, the capacity adjustment cost increases significantly with respect to the number of capacity adjustments, and resultantly the firm has only a single chance to adjust capacity (*a single-adjustment scenario*). This case, in a stylistic way, represents a business environment where it is difficult to change capacity because of managerial hurdles, as illustrated in the earlier examples in bio-pharmaceutical and cosmetics industries. In the second setting, the capacity adjustment cost does not vary with respect to the number of capacity adjustments, and therefore the firm has multiple opportunities to change its capacity (*a multiple-adjustment scenario*) as long as such change results in net benefit. We specifically choose these settings while acknowledging the fact that most decisions in practice fall between these two extremes. We will show that the resultant number of opportunities that the firm has to adjust its capacity critically affects the structure of the optimal policy and asymptotically optimal heuristics.

In both single-adjustment and multiple-adjustment scenarios, we first articulate the stochastic dynamic program formulation of the problem as a special case of the general model and characterize the structure of the optimal policy. Our technical results show that, while optimal polices and derivations are methodologically interesting, they are difficult to implement in practice. For this, we develop data-driven heuristics that are not only implementable but also *asymptotically optimal* with guaranteed convergence rates.

In the single adjustment scenario, we show whether to adjust capacity or not in a given period is not monotone in the firm's posterior belief about demand. In fact, it can be optimal for a firm to increase the capacity level when the likelihood of high demand is moderate, but to stay put (and collect more demand observations) when the likelihood of high demand becomes even higher. Thus, the firm's belief about a demand type does not monotonically affect the decision to adjust capacity or not, while the amount of capacity to adjust is monotone in the firm's belief conditioning on that the firm decides to adjust capacity. In addition, the dynamic program has a very large state space as a state has to include the firm's belief about demand type represented as an information vector. Consequently, solving and implementing the optimal policy quickly becomes computationally intractable even when there are only a few possible demand types. To overcome this, we propose a two-step data-driven heuristic, which only depends on the firm's observed demand data. We prove that this heuristic is *asymptotically optimal* in the case where the true demand follows a stochastic process with stationary and independent increment in time. Specifically, we characterize the rate at which the regret (the percentage profit loss relative to an upper bound when the firm has complete demand information) converges to 0 as the problem scale increases.

We then consider the multiple-adjustment scenario and show that the optimal policy is a control band policy (characterized in Eberly and Van Mieghem 1997), where in each period for a given information vector, the firm will adjust the capacity up to a threshold if the capacity level is significantly low relative to the inferred demand, adjust the capacity down to another threshold if the capacity level is significantly high, and stay put in between. We then further show how the optimal policy changes with respect to the information vector and provide a condition in which the optimal policy changes monotonically. Our results illustrate that computing and implementing the optimal policy are extremely difficult. In place of this, we propose a data-driven heuristic in which the firm adjusts capacity in exponentially increasing intervals and prove that this policy is indeed asymptotically optimal with a provable convergence rate.

The impact of limited capacity adjustment opportunities is highlighted when we compare the optimal policies under the single and multiple adjustment cases. In the single adjustment case, we show that the optimal policy (and the resultant capacity level) is not monotone in the firm's belief about demand being high while in the multiple adjustment case the optimal policy is indeed monotone in the firm's belief about demand being high. In the single adjustment, the firm needs to decide two things: *when to adjust* and *how much*. The non-monotonicity is caused by the option value of adjustment opportunity (if we change the capacity in this period, we cannot change the capacity again). In the multiple adjustment case, however, the firm does not need to worry about exhausting the opportunities (although whether to change or not, and to what level are still driven by capacity cost and leadtime). Although there is a stay-put interval, this is purely driven by cost.

Finally, we illustrate the performance of our heuristics using a numerical study where some of the key parameters and data are derived from actual production and sales data of an automobile instead of using a randomized test bed in order to highlight the fact that our heuristic can be implemented with demand data and a few parameters that can be either inferred or collected by the firm. The numerical study demonstrates the value of using demand learning in capacity decision, and show that our heuristics are very robust with respect to problem parameters and assumptions.

The rest of the paper is organized as follows. The related literature is reviewed in Section 2. The optimal policy of the stochastic dynamic program for the single-adjustment scenario is presented and discussed in Section 4. In Section 4.2, we propose a two-step heuristic and prove its asymptotic optimality. In Section 5 we consider the multiple-adjustment scenario. Similarly to the single-adjustment scenario, we first characterize the optimal policy of a corresponding stochastic dynamic program. We then propose a data-drive heuristic policy under which the firm adjusts its capacity in exponentially increasing intervals and show that this policy is asymptotically optimal. We present the set-up and results of our numerical study in Section 6 and conclude the paper in Section 7.

2. Literature Review

There is an extensive body of literature in the general area of capacity management. Manne (1967), Freidenfelds (1981) and Luss (1982) provide surveys on the earlier literature. In the early work, the main focus is to expand capacity to meet growing demand with no uncertainties. Therefore, the firm is able to make optimal capacity expansion plans to balance economy-of-scale savings and the cost associated with a mismatch between demand and supply. For problems with uncertain demand, Davis et al. (1987) use the piecewise-deterministic Markov process to model an optimal capacity expansion problem with leadtime. Dixit and Pindyck (1994) provide a survey about the real options approach to analyze investment without detailed operational implications. Van Mieghem (2003) provides a comprehensive review about recent developments.

In their seminal paper, Eberly and Van Mieghem (1997) consider a capacity investment problem and present the optimal capacity policy as a control band policy with respect to initial capacity, labeled as the ISD (invest—stay put—divest) policy, when the dynamic capacity adjustment is costly and partially irreversible. The major difference between our work and Eberly and Van Mieghem (1997) is that we explicitly model partially observable states on demand types (which we call information vector) and learning. In particular, we present a concrete scenario on how the evolution of information vector (the probability space) occurs through demand learning and information updating, and study how the *information vector* affects the optimal policy. For instance, in the multiple adjustment setting, we show that the thresholds defining the ISD policy are monotone in the information vector when the information vector is updated according to Bayes' rule. We show that, as the firm's belief of demand being high increases, the firm's decision to adjust capacity changes monotonically with respect to the information vector. We also show, through our analysis with the single adjustment case, that this intuitive result does not hold when the firm has very limited opportunities to adjust capacity. Methodologically, we also extend the result in Eberly and Van Mieghem (1997) to the case where the capacity level is discrete using the $L^\natural$-concave function and its Lovász extension, and establish the optimal policy of the ISD type. We also derive the optimal policy when fixed costs are incurred to adjust capacity. Finally, we propose simple data-driven heuristics with provable analytical bounds.

Among more recent literature on capacity management, a number of papers assume the firm has complete information about the parameterized demand distribution. Among them, Chao et al. (2009) characterize a firm's optimal capacity policy when the existing capacity is subject to deterioration and random supply constraints. Besanko et al. (2010) study an oligopoly in which firms make lumpy capacity investment and disinvestment, and show that while firms build excess capacity for a preemption race in the short run, capacity coordination can be achieved in the long run. Wang et al. (2013) show the optimal capacity policy for two competing technologies is a control band policy. In contrast to these works, our work emphasizes the firm's active role of learning about demand and using it for capacity decisions.

A number of papers consider demand learning in operations contexts. Boyacı and Özer (2010) consider a firm acquiring information via pricing and advance selling, and characterize the firm's optimal policy to stop collecting information and building capacity as a control band policy. Kwon and Lippman (2011) analyze a firm's optimal strategy to invest in project-specific assets with a real option approach, where the firm's profit follows a Brownian motion, and characterize the optimal policy as a control band policy. Kaminsky and Yuen (2014) show a pharmaceutical firm's investment strategy to acquire clinical trial information and build capacity as a threshold policy.

In our paper, we characterize the firm's optimal policy to adjust capacity (increasing or decreasing) in two separate settings – single and multiple adjustment cases. We show that, limiting adjustment opportunities can significantly change the optimal policy and methodology that enables us to characterize the optimal policy. As computing and implementing the optimal solution is extremely difficult when the firm has uncertainties about underlying demand process, we propose a simple and implementable heuristic that is asymptotically optimal with a provable convergence rate. These heuristics overcome challenges posed by incomplete information (Lovejoy 1993).

Our analysis of the optimal policy is closely related to literature on partially observed Markov decision processes (POMDPs), with a particular emphasis on demand learning with Bayesian updating. That is, decision makers know the family of distributions, and update their knowledge about key parameters characterizing the distribution with new observations. Monahan (1982) and Lovejoy (1991) provide surveys about early works in POMDP. Demand learning in a Bayesian fashion has been applied in inventory management (e.g., Scarf 1959, Azoury 1985, Eppen and Iyer

1997, Lariviere and Porteus 1999, Burnetas and Gilbert 2001, Chen and Plambeck 2008). Recently, Aviv and Pazgal (2005) analyze a firm's pricing decision using the POMDP framework. In our paper, we analyze a different operational decision, capacity. Compared to these papers, there are similarities and differences between capacity and inventory decisions. Both capacity and inventory are similar as they are used to satisfy demand, and once acquired, they are costly to maintain (capacity overhead cost and inventory holding cost). Thus, capacity and inventory problems often use the same machinery to derive analytic results. For instance, convexity (concavity) and its preservation over the dynamic program is behind the optimality of the base-stock inventory policy and the optimality of the ISD capacity policy. However, there are several differences between capacity and inventory literatures. First, capacity is a resource used to produce the inventory to satisfy demand. Thus, unless adjustment is made, the capacity level stays the same (regardless of the demand) for each period. On the other hand, inventory is being consumed to meet demand. Thus, the state and its transition in the corresponding dynamic programming formulation are different. In addition, increasing and decreasing capacity are being considered in capacity literature while much of inventory literature (although there are exceptions) does not consider disposal as an active decision. Finally, compared to inventory, capacity adjustment often incurs more significant costs and are subject to managerial hurdles (e.g., the biopharmaceutical plant and the cosmetics company). Therefore, compared to inventory decisions where the firm can reorder every period, the firm may have fewer opportunities to adjust its capacity.

In our work, we explicitly consider the different characteristics of capacity above, formulate the demand learning process of the firm, and analyze both the optimal policy and simple data-driven heuristics. Due to these characteristics of capacity, in our multiple adjustment case, the decision to adjust capacity can be triggered by two or more thresholds with respect to initial capacity while a typical inventory reordering decision is triggered by a single threshold with respect to initial inventory (such as in the base-stock policy or (s, S) policy). In addition, we show through our analysis for the single adjustment case that the limited opportunities to adjust capacity significantly change the firm's optimal policy: although the firm's optimal capacity adjustment decision is monotone with respect to the information vector in the multiple adjustment case, when the opportunities to adjust capacity are very limited, this intuitive result does not hold.

Methodologically, our heuristics are closely related with the recent research on data-driven optimization. Most papers have focused on inventory (Huh et al. 2011, Besbes and Muharremoglu 2013) and pricing (Burnetas and Smith 2000). Some papers also use regret to quantify the heuristics such as Huh and Rusmevichientong (2009) and Besbes and Zeevi (2009). To the best of our knowledge, we are one of the first to apply data-driven optimization in the capacity management setting. In contrast with inventory and pricing decisions, a firm usually has limited opportunities to adjust its capacity, and the adjustment process is often costly and lengthy, which makes the problem somewhat challenging.

3. General Model

In each period, the firm decides whether to change its capacity with existing information and then observe the demand. The firm serves a single product for a finite horizon of J periods, with period 1 and J as the starting and ending periods respectively. The firm operates in a competitive market and is a price-taker so that it cannot adjust price instead of capacity. We assume that each period is of length τ units of time, which will be useful to derive the heuristic in Section 4.2.

We assume that the firm has incomplete information about the demand: while the firm knows the demand pattern or distribution family about the demand, some key parameters characterizing the demand (which we call *demand type*) are unknown. Specifically, there are $I \in \mathbb{N}$ potential demand types: θ_i for $i \in \{1, 2, ..., I\}$ and $\theta_{i_1} < \theta_{i_2}$ if $i_1 < i_2$. The parameter, θ_i, determines the parameter(s) of the underlying demand distribution. Thus, for given demand type i, the demand in period j, D_j, is represented by a random variable $D_j|\theta_i = \lambda_j(\theta_i) + \xi_j|\theta_i$ where $\lambda_j(\theta_i)$ is the mean demand of $D_j|\theta_i$, and $\xi_j|\theta_i$ is a random term with mean 0. We assume the random term ξ_j is independent across periods. A number of demand processes can be expressed in this way and our results on the optimal policy apply to a large class of random variables and demand processes (see remark on demand process in Section EC.2.1.2).

We assume that demand in each period is stochastically ordered in the demand type parameter: $D_j|\theta_{i_1} \preceq_{st} D_j|\theta_{i_2}$ for $i_1 \leq i_2$. Thus, demand stochastically increases in the demand type index, i. We use $F_j(\cdot|\theta_i)$ and $f_j(\cdot|\theta_i)$ to denote the cumulative distribution and the density function (probability mass function in the case of discrete demand) of $D_j|\theta_i$. Finally, for ease of exposition, we write $\lambda_j(\theta_i)$ as $\lambda_{j,i}$, and assume $\lambda_{j,I} < \infty$ for all j for analytical tractability.

The firm observes demand and uses the observations to update its belief about the true demand type. The firm's information about the demand evolves as follows. Let the vector $\boldsymbol{\pi_1}$ be the firm's prior distribution of the demand type at the beginning of period 1: $\boldsymbol{\pi_1} = (\pi_{1,1}, \ldots, \pi_{1,I})$ where $\pi_{1,i} = \Pr(\Theta = \theta_i)$. At the beginning of period j ($j > 1$), the firm's information about the demand type is represented by an information vector $\boldsymbol{\pi_j} \triangleq (\pi_{j,1}, \pi_{j,2}, ..., \pi_{j,I})$. The $\pi_{j,i}$ is defined as the posterior distribution of the demand being type i given the past demand history, i.e., $\pi_{j,i} \triangleq \Pr(\Theta = \theta_i|\boldsymbol{d_{j-1}})$ where $\boldsymbol{d_{j-1}} \triangleq (d_1, d_2, ..., d_{j-1})$ is a demand history up to period $j-1$. After the firm observes d_j, demand at the end of period j, the information vector is updated following Bayes' rule:

$$\pi_{j+1,i} = \frac{\pi_{j,i} f_j(d_j|\theta_i)}{\sum_{k=1}^{I} [\pi_{j,k} f_j(d_j|\theta_k)]}. \tag{1}$$

Before the realization of D_j, the information vector is a vector of random variables (denoted by $\mathbf{\Pi_{j+1}}$). The following lemma, adapted from page 96 of Williams (1991), proves that given the current distribution about the demand types, the conditional posterior distributions in the future periods are the same as the current one in expectation. Thus, the posterior distribution has the martingale property.

LEMMA 1 **(Martingale property of the posterior distribution, Williams 1991)**.

$$E\left[\mathbf{\Pi_{j_2}}|\mathbf{\Pi_{j_1}}\right]=\mathbf{\Pi_{j_1}},\ for\ j_1\le j_2. \tag{2}$$

In each period, the firm observes the realized demand, d_j, and fulfills the demand using the firm's existing capacity in that period. For each unit it satisfies with existing capacity, the firm accrues a profit of p, which represents the revenue minus the variable production cost (excluding any capacity cost). If demand exceeds the firm's capacity, we assume that it is satisfied by an outside option such as overtime production, emergency fulfillment, or using other production facilities. Let c_1 be the per-unit outside option cost (we call this *outside option* or *outsourcing* cost). Note that c_1 represents the cost premium of producing one unit using the firm's outside option. In addition to production costs, the firm also incurs an overhead cost to maintain the existing capacity, denoted by c_0 per unit capacity and unit time. As this cost represents the firm's cost to own and maintain the capacity, it is incurred whether the capacity is used or not in that period. To avoid trivial cases, we assume $p\ge c_1>c_0$, i.e., the unit profit is higher than the unit cost associated with the outside option, otherwise the firm will not outsource any demand; the unit outsourcing cost is higher than the cost to maintain one unit of the firm's own capacity for one period, otherwise the firm will not have incentive to build any capacity. A similar cost structure was used in Chao et al. (2009).

When the firm's capacity level is μ and the firm's belief about the demand type is $\boldsymbol{\pi_j}$, the firm's expected operating profit $h_j(\boldsymbol{\pi_j},\mu)$ in period j (note that each period is τ units of time) is as follows, where $x^+\triangleq\max\{x,0\}$

$$\begin{aligned}h_j(\boldsymbol{\pi_j},\mu)&\triangleq E_\Theta\left\{E_{D_j|\Theta}\left[pD_j-c_1\left(D_j-\mu\tau\right)^+-c_0\mu\tau|\Theta\right]\Big|\,\boldsymbol{\pi_j}\right\}\\&=\sum_{k=1}^{I}\pi_{j,k}E\left[pD_j-c_1\left(D_j-\mu\tau\right)^+-c_0\mu\tau|\Theta=\theta_k\right].\end{aligned}\tag{3}$$

We note that our base model has a number of assumptions such as (i) demand is not censored, (ii) leadtimes are symmetric (leadtime for building capacity is the same as leadtime for disposing capacity), and (iii) only variable capacity costs are assumed. We show that relaxing these assumptions does not change our results and analysis: see the remarks in Section EC.2.1.2.

We next describe the firm's capacity decision. At the beginning of the planning horizon, the firm has initial capacity, μ_0. This is for generality. The firm may start with no existing capacity: $\mu_0=0$, or the firm may use the prior distribution to choose a capacity level or use the existing (legacy) capacity: $\mu_0 > 0$. In each period, the firm decides whether it should adjust its capacity. As changing capacity often requires considerable amount of time, we assume there is a leadtime of l periods.

To be more specific, suppose that the firm has a capacity level of μ in period j.[1] If the firm decides to change the capacity level from μ to μ' in period j, the firm's existing capacity will be changed to μ' after l periods (in period $j+l$). We assume that both increasing and decreasing the capacity level are costly to the firm. As we noted earlier, the capacity adjustment cost may depend on the number of adjustment the firm has made. Let η denote the number of capacity adjustments that firm has already made, and $c_{a,\eta}$ be the cost of adding one unit of capacity and $\gamma_{a,\eta}$ be the cost of decreasing one unit of capacity. Thus, the cost associated with changing the capacity level from μ to μ', denoted by $\hat{C}(\mu,\mu';\eta)$, is

$$\hat{C}(\mu,\mu';\eta) \triangleq c_{a,\eta}(\mu'-\mu)^+ + \gamma_{a,\eta}(\mu-\mu')^+, \text{ where } c_{a,\eta} = c_a + \Lambda\eta \text{ and } \gamma_{a,\eta} = \gamma_a + \Lambda\eta. \tag{4}$$

Notice that if the firm does not change the capacity, $\hat{C}(\mu,\mu;\eta) = 0$. We assume $c_{a,\eta} \geq 0$ and $c_{a,\eta} + \gamma_{a,\eta} \geq 0$, indicating that it is costly to reverse the installed capacity. Following our discussions earlier, we assume that $c_{a,\eta}$ and $\gamma_{a,\eta}$ weakly increase in η, i.e., $\Lambda \geq 0$, which means that the more capacity adjustment the firm commits to, the more expensive it becomes to adjust capacity. The parameter Λ measures the degree of managerial hurdles in capacity adjustment. The greater the hurdle is, the larger the Λ is. For example, Λ is larger in the case of the cosmetics production manager than the case of the data center. Note that $\gamma_{a,\eta} < 0$ implies that the firm may salvage a portion of its capacity cost, and $\gamma_{a,\eta} \geq 0$ implies that downsizing the capacity is costly to the firm (e.g., cost of layoff). To avoid trivial cases, we also assume $c_1(J-l)\tau \geq c_{a,0} + c_0(J-l)\tau$ and $c_0(J-l)\tau \geq \gamma_{a,0}$. The first assumption implies that it is less costly to increase a unit of capacity for the first time and maintain it than outsourcing this unit to the more expensive outside option for the whole time after the adjustment. The second assumption implies that it is cheaper to dispose unnecessary capacity for the first time than holding it for the whole time after the adjustment. To simplify notations, we define $\hat{\mu}_j$ as the capacity position in period j. In general, if the capacity position is $\hat{\mu}_j$ in period j, the actual capacity level in period $j+l$ is $\mu_{j+l} = \hat{\mu}_j$. In the example above, if the firm decides to adjust capacity from μ to μ' in period j, the capacity position $\hat{\mu}_j$ is changed to μ' immediately while the actual capacity level will be updated to μ' after l periods.

We allow the set of capacity levels (denoted by $\mathcal{K}$) to be discrete or continuous. When capacity level is primarily determined by the number of key machines or production lines, it may be appropriate that the capacity level must be chosen from a discrete set, i.e., $\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$, where $\mathbb{Z}_+$ stands for the set of nonnegative integers. Otherwise, capacity levels can be continuous (e.g., the capacity is measured by the available labor hours), i.e., $\mathcal{K} = \mathbb{R}_+$, where $\mathbb{R}_+$ stands for the set of nonnegative real numbers.

Given state $(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j)$, we next derive $V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j)$, the firm's value-to-go function starting at period j. We first define $H_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}_j, \eta_j)$ to be the expected operating profit in period $j+l$ (that is, when the capacity is in effect) minus the capacity adjustment cost that the firm incurs, provided that the firm has adjusted the capacity η_j times.

$$H_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}_j, \eta_j) \triangleq E\left[h_{j+l}(\boldsymbol{\Pi_{j+l}}, \hat{\mu}_j) - \hat{C}(\hat{\mu}_{j-1}, \hat{\mu}_j; \eta_j) | \boldsymbol{\pi_j}\right] = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}_j) - \hat{C}(\hat{\mu}_{j-1}, \hat{\mu}_j; \eta_j) \quad (5)$$

The equality follows Lemma 1 and the fact that $h_{j+l}(\boldsymbol{\Pi_{j+l}}, \hat{\mu}_j)$ is linear in $\boldsymbol{\Pi_{j+l}}$. We also define a policy g as $\{\hat{\mu}_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j), j = 1, 2, ..., J-l\}$ and $\mathcal{G}^m$ as the set of all the admissible policies. Then, the firm's problem is to determine a policy $g^* \in \mathcal{G}^m$ to maximize the total expected profit,

$$\max_{g \in \mathcal{G}^m} \sum_{j=1}^{l} E\left[h_j(\boldsymbol{\Pi_j}, \mu_0) | \boldsymbol{\pi_1}\right] + \sum_{j=1}^{J-l} E^g\left[H_j(\boldsymbol{\Pi}_j, \hat{\mu}_{j-1}, \hat{\mu}_j, \eta_j) | \boldsymbol{\pi_1}\right] \quad (6)$$

where the expectation is taken over D_j for all j at time zero. As the profit from the first l periods is not affected by the firm's capacity decision, the decision problem is to find a policy that maximizes the following function:

$$V_1(\boldsymbol{\pi_1}, \hat{\mu}_0, 0) = \max_{g \in \mathcal{G}^m} \sum_{j=1}^{J-l} E^g\left[H_j(\boldsymbol{\Pi}_j, \hat{\mu}_{j-1}, \hat{\mu}_j, \eta_j) | \boldsymbol{\pi_1}\right] \quad (7)$$

We define an indicator $\mathbf{1}_C$: $\mathbf{1}_C = 1$ if condition C is met and 0 otherwise. Then the optimal value-to-go function is recursively defined as follows: for all $j \in \{1, 2, ..., J-l\}$,

$$\begin{aligned} V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) = \max_{\hat{\mu} \in \mathcal{K}} E\Big[& H_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}, \eta_j) + V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}, \eta_j) \mathbf{1}_{\{\hat{\mu} = \hat{\mu}_{j-1}\}} \\ & + V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}, \eta_j + 1) \mathbf{1}_{\{\hat{\mu} \neq \hat{\mu}_{j-1}\}} | \boldsymbol{\pi_j} \Big] \\ V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) = 0 \text{ for } & j > J - l. \end{aligned} \quad (8)$$

Equivalently, by defining a state-dependent fixed cost $K(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) \triangleq E[V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1}, \eta_j) - V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1}, \eta_j + 1) | \boldsymbol{\pi_j}]$, we have equations (8) reformulated as follows:

$$\begin{aligned} V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) = \max_{\hat{\mu} \in \mathcal{K}} \Big\{ & h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - \hat{C}(\hat{\mu}_{j-1}, \hat{\mu}; \eta_j) - K(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) \mathbf{1}_{\{\hat{\mu} \neq \hat{\mu}_{j-1}\}} \\ & + E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}, \eta_j + 1) | \boldsymbol{\pi_j}\right] \Big\} + K(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j); \\ V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j) = 0 \text{ for } & j > J - l. \end{aligned} \quad (9)$$

Therefore, the decision problem is equivalent to one featuring both demand learning as reflected by the state $\boldsymbol{\pi_j}$ and a state-dependent capacity adjustment *fixed cost* $K(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \eta_j)$. The state-dependent fixed cost makes the problem generally intractable. In fact, we show that even for a capacity investment problem with demand learning and a stationary fixed cost, the optimal policy can be rather complicated in Section EC.2.2.2.

In order to clearly understand the impact of demand learning on the capacity investment decision, we consider two special cases in the following sections. Recall that $c_{a,\eta} = c_a + \Lambda\eta$ and $\gamma_{a,\eta} = \gamma_a + \Lambda\eta$ where Λ measures the degree of managerial hurdles. We first consider the case where the capacity adjustment cost increases significantly with the number of adjustments, i.e., $\Lambda = M$ where M is a very large number, reflecting the scenario with significant managerial hurdles such as the case for the cosmetic production manager. As a result, the firm has only one opportunity to adjust its capacity over the decision horizon. We analyze this case in Section 4. We also consider the case where the capacity adjustment cost does not change with respect to the number of adjustments, i.e., $\Lambda = 0$, reflecting the scenario of little managerial hurdles in capacity adjustment. In this scenario, the firm has multiple opportunities to adjust its capacity. We analyze this case in Section 5.

4. Capacity Investment with a Single Adjustment Opportunity

We first study the case where the firm's capacity adjustment cost increases significantly in the number of adjustments, i.e., the managerial hurdle parameter $\Lambda = M$ where M is a very large number. In this case, it follows that the firm has a single opportunity to adjust capacity during a planning horizon. The state of initial capacity μ_0 is suppressed for simplicity and the number of capacity adjustment is embedded in the new formulation. Therefore, we only need to show the state of the firm's information on the demand type as the state.

We derive $V_j(\boldsymbol{\pi_j})$, the firm's value-to-go function starting at period j, as follows. If the firm has not adjusted the capacity before, then the firm can choose either to adjust capacity, whose value corresponds to $L_j^a(\boldsymbol{\pi_j})$, or to stay put in period j, whose value corresponds to $L_j^s(\boldsymbol{\pi_j})$. Formally, the value-to-go function $V_j(\boldsymbol{\pi_j})$ is defined as follows[2]:

$$V_j(\boldsymbol{\pi_j}) = \max\left\{L_j^a(\boldsymbol{\pi_j}), L_j^s(\boldsymbol{\pi_j})\right\} \tag{10}$$

If the firm chooses to adjust capacity in the period, then the firm chooses the capacity that maximizes the expected total profit. Thus, $L_j^a(\boldsymbol{\pi_j})$, where the superscript a is for "adjustment", is:

$$\begin{aligned} L_j^a\left(\boldsymbol{\pi_j}\right) &\triangleq \max_{\mu\in\mathcal{K}} E\left[\sum_{k=j+l}^{J} h_k(\boldsymbol{\Pi_k},\mu) - \hat{C}(\mu_0,\mu)\middle|\, \boldsymbol{\pi_j}\right] = \max_{\mu\in\mathcal{K}}\left\{\sum_{k=j+l}^{J} h_k(\boldsymbol{\pi_j},\mu) - \hat{C}(\mu_0,\mu)\right\} \\ &= \sum_{k=j+l}^{J} h_k\left(\boldsymbol{\pi_j}, \hat{\mu}_j^a(\boldsymbol{\pi_j})\right) - \hat{C}\left(\mu_0, \hat{\mu}_j^a(\boldsymbol{\pi_j})\right) \end{aligned} \tag{11}$$

The first equality follows Lemma 1 and the fact that $h_k(\boldsymbol{\Pi_k},\mu)$ is linear in $\boldsymbol{\Pi_k}$. We also define the induced target capacity position $\hat{\mu}_j^a(\boldsymbol{\pi_j}) = \arg\max_{\mu\in\mathcal{K}}\left\{\sum_{k=j+l}^{J} h_k(\boldsymbol{\pi_j},\mu) - \hat{C}(\mu_0,\mu)\right\}$, the optimal capacity position to adjust in period j for given information vector $\boldsymbol{\pi_j}$ if the firm decides to adjust capacity. Later in the paper, for simplicity, we write it as $\hat{\mu}^a_j$ when there is no confusion. In case there are multiple maximizers, we break the tie by choosing the smallest capacity position.

If the firm stays put in period j, the value-to-go function $L^s_j(\boldsymbol{\pi_j})$, where we use the superscript s for "stay put", is defined as

$$L^s_j(\boldsymbol{\pi_j}) \triangleq h_{j+l}(\boldsymbol{\pi_j}, \mu_0) + E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}})\right| \boldsymbol{\pi_j}] \tag{12}$$

which is the sum of the two terms: the expected one period profit when capacity is μ and the expected future profit (from period $j+1$). Note that, if the firm stays put in period j, it maintains the option to change the capacity in the future.

Therefore, to characterize the firm's optimal policy, we only need to compare $L^a_j(\boldsymbol{\pi_j})$ and $L^s_j(\boldsymbol{\pi_j})$: the choice between making an adjustment in period j or delaying the decision. Note that, in the single adjustment case, the problem of choosing "when to adjust" and "by how much" is recast as an optimal stopping time problem.

4.1. Optimal Policy in the single adjustment case

We now characterize the firm's optimal capacity policy, starting with the case when capacity levels form a discrete set ($\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$). To this end, we first define a convex partition of the set that contains feasible information vectors $\boldsymbol{\pi_j}$: $\mathcal{P}_j = \{\boldsymbol{\pi_j} = (\pi_{j,1}, \pi_{j,2}, \ldots, \pi_{j,I}) : \sum_{i=1}^{I} \pi_{j,i} = 1, \pi_{j,i} \geq 0\}$.

Definition 1. $\mathbb{P}_j = \{\mathcal{P}_{j,k} : \mathcal{P}_{j,k} \subset \mathcal{P}_j\}$ is a convex partition of $\mathcal{P}_j$, if the following conditions are satisfied:

(i) $\emptyset \notin \mathbb{P}_j$;

(ii) $\bigcup_k \mathcal{P}_{j,k} = \mathcal{P}_j$;

(iii) if $k \neq r$, then $\mathcal{P}_{j,k} \bigcap \mathcal{P}_{j,r} = \emptyset$;

(iv) for any $\alpha \in (0,1)$, if $\boldsymbol{\pi_j} \in \mathcal{P}_{j,k}$ and $\boldsymbol{\hat{\pi}_j} \in \mathcal{P}_{j,k}$, then $\alpha\boldsymbol{\pi_j} + (1-\alpha)\boldsymbol{\hat{\pi}_j} \in \mathcal{P}_{j,k}$.

In other words, $\mathbb{P}_j$ is a collection of subsets of information vectors where each subset is non-empty and convex, and the union of these subsets is $\mathcal{P}_j$.

We next characterize the firm's optimal policy to adjust the capacity. We use $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi'_j}$ to denote that the posterior distribution $\boldsymbol{\pi_j}$ is smaller than $\boldsymbol{\pi'_j}$ in the first order stochastic sense, i.e., $\sum_{k=1}^{i} \pi_{j,k} \geq \sum_{k=1}^{i} \pi'_{j,k}$ for all $i = 1, 2, ..., I$. Let $\hat{\mu}^*_j(\boldsymbol{\pi_j})$ be the optimal capacity position in period j given information vector $\boldsymbol{\pi_j}$, $\mathcal{S}_{j,k}$ be a sequence of convex sets, and $S_j = \bigcup_k S_{j,k}$. The following proposition characterizes the optimal policy.

Proposition 1 **(Optimal capacity policy: Discrete capacity case)**. *For $j = 1, 2, ..., J-l$:*

(i) $L^a_j(\boldsymbol{\pi_j})$ and $L^s_j(\boldsymbol{\pi_j})$ are convex in $\boldsymbol{\pi_j}$. Therefore, $V_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$.

(ii) Let $\mathcal{P}_{j,k} = \left\{\boldsymbol{\pi_j} : \hat{\mu}^a_j(\boldsymbol{\pi_j}) = k\delta\right\}$. Then, $\mathbb{P}_j = \{\mathcal{P}_{j,k} : \mathcal{P}_{j,k} \neq \emptyset, k \in \mathbb{Z}_+\}$ is a convex partition of $\mathcal{P}_j$.

(iii) In each $\mathcal{P}_{j,k} \in \mathbb{P}_j$*, there exists at most one convex set* $\mathcal{S}_{j,k} \subseteq \mathcal{P}_{j,k}$ *such that if* $\boldsymbol{\pi_j} \in \mathcal{S}_{j,k}$*, it is optimal to adjust the capacity position to* $k\delta$*:* $\hat{\mu}_j^*(\boldsymbol{\pi_j}) = \hat{\mu}_j^a(\boldsymbol{\pi_j}) = k\delta$*. If* $\boldsymbol{\pi_j} \notin S_j$*, then it is optimal to wait:* $\hat{\mu}_j^*(\boldsymbol{\pi_j}) = \mu_0$*.*
(iv) If $\boldsymbol{\pi_j}, \boldsymbol{\pi_j'} \in S_j$ *and* $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi_j'}$*, then* $\hat{\mu}_j^*(\boldsymbol{\pi_j}) \leq \hat{\mu}_j^*(\boldsymbol{\pi_j'})$*.*

Part (iii) of Proposition 1 implies that the firm's decision to adjust capacity level in the current period is not monotone in its belief about the demand type, $\boldsymbol{\pi_j}$. That is, it is possible that the firm may increase the capacity when the likelihood of high demand is low, but wait to observe more demand when the likelihood of high demand becomes even higher, i.e., $\hat{\mu}_j^*(\boldsymbol{\pi_j}) > \hat{\mu}_j^*(\boldsymbol{\pi_j'})$ for $\boldsymbol{\pi_j} \prec \boldsymbol{\pi_j'}$. Thus, as $\boldsymbol{\pi_j}$ stochastically increases, the optimal policy can switch multiple times between waiting and adjusting. Within each $\mathcal{P}_{j,k} \in \mathbb{P}_j$, it is optimal to adjust capacity only when $\boldsymbol{\pi_j}$ falls in a convex subset $\mathcal{S}_{j,k}$. If $\boldsymbol{\pi_j} \in \mathcal{P}_{j,k} \backslash \mathcal{S}_{j,k}$, it is optimal to wait. Recall that the problem defined in equation (10) is indeed an optimal stopping problem. Thus, one would expect that the optimal policy would be characterized by a monotone threshold (switching curve) in information vector as $\boldsymbol{\pi_j}$ stochastically increases. Proposition 1 shows that it is not the case. Although this seems surprising at first, this phenomenon indeed reflects the primary trade-off that the firm juggles—exploration versus exploitation. On one hand, the firm would like to exploit benefits from the current information by adjusting the capacity now. On the other hand, if a few more observations of the demand (and, resultantly, updated belief) may shift the firm's target capacity position considerably, it might be beneficial to wait. Part (iv) shows that in regions where it is optimal to change the capacity position, the target capacity position increases in the information vector, $\boldsymbol{\pi_j}$. In other words, given that the firm changes the capacity in the same period, the optimal capacity position is monotonically increasing in the information vector.

Figure 1 illustrates how the optimal policy changes in $\boldsymbol{\pi_j}$ in a two-demand-type case. We note that as there are only two demand types, the information vector is $\boldsymbol{\pi_j} = (\pi_{j,1}, \pi_{j,2})$ and $\pi_{j,1} + \pi_{j,2} = 1$, so it is sufficient to show how the optimal policy changes with respect to $\pi_{j,2}$, the probability of demand being type high. In this case, the space of feasible information vectors $\mathcal{P}_j$ is partitioned into 5 convex subsets ($\mathcal{P}_{j,k_i}$), and each subset corresponds to a different level of $\hat{\mu}_j^a(\boldsymbol{\pi_j})$ (i.e., the induced target capacity position given the firm decides to adjust capacity in that period shown by the dashed lines in Figure 1(B)). The white regions correspond to the regions in which it is optimal to adjust the capacity ($\mathcal{S}_{j,k}$) while the grey regions correspond to the stay-put region. Notice that the firm may choose to adjust capacity for the whole region ($\mathcal{S}_{j,k_1} = \mathcal{P}_{j,k_1}$), or choose to wait for the whole region ($\mathcal{S}_{j,k_4} = \emptyset$).

In Figure 1(B), we observe that $\hat{\mu}^a$ (indicated by dashed lines) increases in the information vector, i.e., $\hat{\mu_j}^a(\boldsymbol{\pi_j}) \leq \hat{\mu}^a(\boldsymbol{\pi_j'})$ when $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi_j'}$. However, the decision to adjust the capacity is not

Figure 1 **A numerical example of the optimal policy with two demand types and discrete capacity levels.**

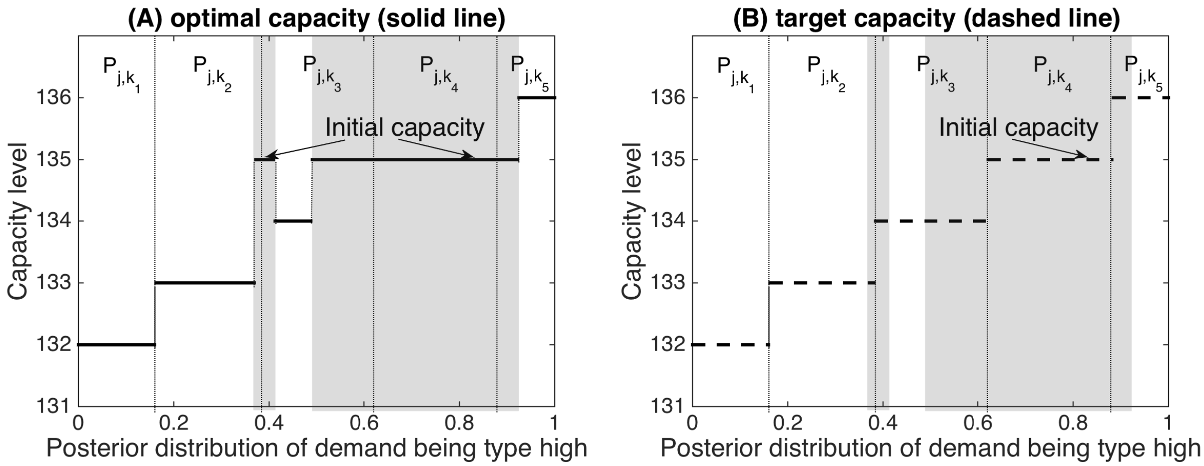

Note. Parameters: price $p = 7$, capacity outsourcing cost $c_1 = 7$, capacity overhead cost $c_0 = 1$, capacity upgrading cost $c_a = 0$, capacity downgrading cost $\gamma_a = 0$; high demand type follows a Poisson distribution with mean of 124; low demand type follows a Poisson distribution with mean of 120. Initial capacity $\mu_0 = 135$. Base capacity unit $\delta = 1$. Leadtime $l = 0$. There are 20 periods in total and the example illustrates the policy in period 12. The white regions correspond to the adjustment regions while the grey regions correspond to the stay-put region. When the optimal decision is to stay put, the optimal capacity corresponds to the initial capacity, as defined in Proposition 1.

monotone in $\boldsymbol{\pi_j}$ as shown in Figure 1(A). In this case, given that the initial capacity corresponds to the optimal target position in region 4, as the information vector increases, the optimal decision on when to change the capacity is not monotone. The firm first chooses to adjust down (regions $\mathcal{P}_{j,k_1}$ to $\mathcal{P}_{j,k_2}$), then stay put (regions $\mathcal{P}_{j,k_2}$ to $\mathcal{P}_{j,k_3}$), then adjust down again (region $\mathcal{P}_{j,k_3}$), then stay put again (regions $\mathcal{P}_{j,k_3}$ to $\mathcal{P}_{j,k_5}$), and finally adjust up (region $\mathcal{P}_{j,k_5}$).

While one may think that this discontinuity is driven by the fact that the feasible capacity level must be chosen from a discrete set, we show that the same result holds even when the capacity level is a continuous variable, as shown in the next proposition. We define $\hat{\mu}^*_j(\boldsymbol{\pi_j})$ as the optimal capacity position in period j given information vector $\boldsymbol{\pi_j}$ and let $S_j = \{\boldsymbol{\pi_j} : L^a_j(\boldsymbol{\pi_j}) > L^s_j(\boldsymbol{\pi_j})\}$.

Proposition 2 **(Optimal capacity policy: Continuous capacity case)**. *For $j = 1, ..., J - l$,*

(i) $L^a_j(\boldsymbol{\pi_j})$ and $L^s_j(\boldsymbol{\pi_j})$ are convex in $\boldsymbol{\pi_j}$. Therefore, $V_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$.

(ii) For $\boldsymbol{\pi_j}$ and $\boldsymbol{\pi'_j} \in S_j$, if $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi'_j}$, then $\hat{\mu}^*_j(\boldsymbol{\pi_j}) \leq \hat{\mu}^*_j(\boldsymbol{\pi'_j})$.

Part (ii) above implies that as in the discrete type case if the firm decides to adjust the capacity, then the target capacity position increases in the information vector. However, it is important to notice that the optimal capacity position under the optimal policy is still non-monotone in the information vector. This is because the firm's decision about whether to adjust the capacity in this period does not change monotonically with respect to the increased likelihood. To see why,

Figure 2 **An example of non-monotone optimal adjustment region across periods.**

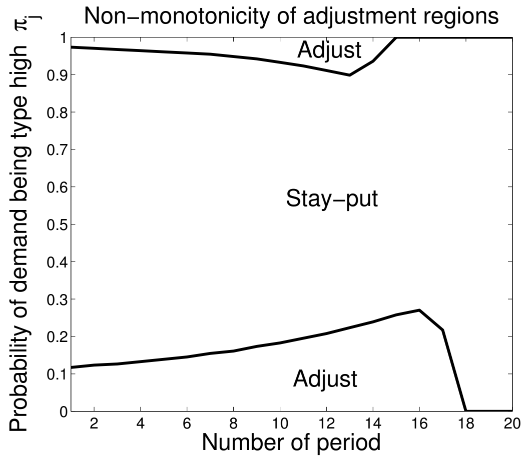

Note. Parameters: price $p = 60$, capacity outsourcing cost $c_1 = 60$, capacity overhead cost $c_0 = 30$, capacity upgrading cost $c_a = 100$, capacity downgrading cost $\gamma_a = 50$; high demand type follows a Gamma distribution with mean of 20 and variance of 30; low demand type follows a Gamma distribution with mean of 10 and variance of 15. Initial capacity $\mu_0 = 25$. Leadtime $l = 0$. There are 20 periods in total. π_j is the posterior probability of demand being type high in period j.

note that when-to-stop decision is determined by comparing two functions that are convex in $\boldsymbol{\pi_j}$, $L_j^a(\boldsymbol{\pi_j})$ and $L_j^s(\boldsymbol{\pi_j})$, in the optimal stopping problem. Thus, S_j is not necessarily a connected set.

So far, we have discussed the firm's optimal policy in any given period. Next, we discuss how the optimal policy changes across periods. We find that, with capacity cost (even when only the variable cost is incurred), the optimal policy is non-monotone in time. To see this, we have first considered the case where it is costless to adjust capacity. In this case, we show that the regions in which the firm changes capacity (either increasing or decreasing) expands in time. This makes sense, as time moves towards the end of the horizon, the value of staying-put (and observing demand) decreases because the time left to utilize the explored knowledge to maximize profit becomes shorter.

As illustrated in Figure 2, this is no longer true if changing capacity is costly, i.e., $c_a > 0$ and $\gamma_a > 0$. In the figure, the adjustment region first expands, and then shrinks in time. As the firm gets enough information, the value of staying-put (and observing demand) decreases because the time left to maximize profit also decreases, which explains the expansion region. As the period approaches the end of the horizon, the return from the adjusted capacity may not justify the cost to adjust it. In this case, the information becomes of almost no value as there is not enough time to recoup investment. Thus, as seen in the figure, the adjustment region shrinks as time moves towards the end of the horizon even when there is no capacity leadtime.

We briefly discuss modeling features and assumptions on demand process, fixed cost and leadtime, the rationale behind them, and the consequences of removing or relaxing them in Appendix

EC.2.1.2. In Section 5, we will consider the multiple adjustment case and highlight the difference. The subsection following the remarks will derive a heuristic policy and analyze its performance.

4.2. Near-Optimal Heuristic and Performance Evaluation

As Section 4.1 shows, the optimal policy is complicated and difficult to implement even for small problems with finite demand types and capacity levels. One of the reasons is that the state space—which includes the information vector $\boldsymbol{\pi_j}$—is uncountably infinite, and therefore computing the exact optimal policy is computationally intractable for large problems. For smaller problems, a fine mesh approximation with linear interpolation can approximate the value-to-go function and hence the optimal policy (as we do for one part of our validation case study in Section 6), but in general, the curse of dimensionality makes it impossible to find the optimal policy.

Therefore, we propose a simple two-step heuristic and prove its performance. The firm observes demand for a specific amount of time (τ_n units of time whose value depends on the problem size) and then adjusts the capacity based on the observed demand. We then show that, in an asymptotic regime, this heuristic is near-optimal when the underlying demand process is a stationary process with unknown mean under the *regret* criterion, which quantifies the gap between an upper-bound (based on information relaxation and deterministic approximation) and the value-to-go function derived from the two-step heuristic.

We scale up both demand and capacity by a coefficient n to define the asymptotic regime. For this, consider the firm's problem with a planning horizon $[0,T]$ within which the firm reviews its decision periodically. Let τ_n be the time between two consecutive decision opportunities so that the corresponding decision problem is a discrete-time dynamic program with $J_n = T/\tau_n$ periods. Likewise, let l_n be the capacity lead time (described in the number of periods): $l_n = l_t/\tau_n$ where $l_t \in [0,T]$. Without loss of generality, we assume $J_n = T/\tau_n$ and $l_n = l_t/\tau_n$ are integers. Note that the choice of τ_n affects the performance of a heuristic. In our heuristic, we construct a sequence of τ_n to achieve the asymptotic optimality with a provable convergence rate. We note that continuous monitoring of the demand process may not be feasible, for example, when the demand information needs to be aggregated from retailers at various locations. In these cases, firms often use a periodic reporting rule so that the realized demand can be collected and reported in a periodic fashion.

We assume that the firm's demand follows a stationary random process with an unknown average demand rate. Let $\{N(t), t \geq 0\}$ denote a standard random process with stationary and independent increment, which satisfies $N(0) = 0$, has mean $E[N(t)] = t$, and variance $Var[N(t)] = \sigma^2 t$ for $t \geq 0$. For example, when $\sigma = 1$, the process $\{N(t), t \geq 0\}$ may represent a unit-rate Poisson process, as the one considered in Besbes and Zeevi (2009); when $\sigma \neq 1$, $\{N(t), t \geq 0\}$ may represent a compound Poisson process. When the demand type is i, we define $\{N(n\lambda_i t), t \geq 0\}$, $i \in \{1, 2, ..., I\}$ as the

Table 1 The two-step heuristic

	Given the period length of τ_n,
1.	The firm serves the demand in period 1 with initial capacity $n\mu_0$. Let $n\hat{\lambda}_{\tau_n}$ be the observed demand rate in period 1.
2.	The firm adjusts its capacity position to $n\hat{\lambda}_{\tau_n}$.
3.	The firm serves demand from period 2 to l_n+1 with the initial capacity $n\mu_0$, and from l_n+2 to J_n with capacity $n\hat{\lambda}_{\tau_n}$.

corresponding demand process with demand type parameter $\theta_{i,n} = n\lambda_i$. That is, given demand type i, the firm's demand in period j is $D_j|\theta_{i,n} = N(n\lambda_i j\tau_n) - N(n\lambda_i(j-1)\tau_n)$, and therefore the demand in each period is a sequence of i.i.d. random variables with mean $n\lambda_i\tau_n$ and variance $\sigma^2 n\lambda_i\tau_n$.[3]

We assume that the firm's initial capacity is scaled up as $n\mu_0$. All other aspects of the model (e.g., costs, revenue, etc.) are the same as the original model considered in Section 4. To show the asymptotic optimality, we impose the following assumptions on λ_i such that $\lambda_i \in [0,\ M]$. The assumption stipulates that, the demand rate for any type i is bounded from above. With this set-up, we now introduce and analyze the two-step heuristic, denoted by (ts).

The two-step heuristic. The firm observes demand for one period comprising τ_n units of time, and then uses the observed demand rate to adjust the capacity for the rest of the time horizon, as specified in Table 1. We will show that this simple policy is asymptotically optimal with an appropriately chosen τ_n (which is a function of the scale parameter n).

Under the two-step heuristic, the firm always adjusts the capacity to the observed demand rate in the first period. Therefore, we define

$$\hat{\lambda}_{i,\tau_n} \triangleq \frac{N(n\lambda_i\tau_n)}{n\tau_n}, \tag{13}$$

Then the firm's expected value-to-go function under the heuristic is as follows.

$$V_{0,n}^{ts}(\boldsymbol{\pi_1}) = \sum_{i=1}^{I} \pi_{1,i} E \left\{ \begin{array}{l} pD_{l_n+1} - c_1\left(D_{l_n+1} - n\mu_0\tau_n\right)^+ - c_0 n\mu_0\tau_n - \hat{C}(n\mu_0, n\hat{\lambda}_{i,\tau_n}) \\ + \sum_{j=l_n+2}^{J_n} \left[pD_j - c_1\left(D_j - n\hat{\lambda}_{i,\tau_n}\tau_n\right)^+ - c_0 n\hat{\lambda}_{i,\tau_n}\tau_n \right] \end{array} \middle| \theta_{i,n} \right\} \tag{14}$$

For ease of exposition, we suppress the dependency of $V_{0,n}^{ts}(\boldsymbol{\pi_1})$ on $\boldsymbol{\pi_1}$ when there is no confusion.

As the two-step heuristic is a feasible policy for the corresponding optimal stopping problem, it follows that the value-to-go function under the two-step heuristic, $V_{0,n}^{ts}$, is a lower bound of the value-to-go function under the optimal policy, denoted by $V_{0,n}^{*}$. However, because of the complexity of the optimal policy and the curse of dimensionality, the exact value function under the optimal policy, denoted by $V_{0,n}^{*}$ is difficult to compute. Hence, we will introduce an upper-bound of $V_{0,n}^{*}$ to evaluate the performance of the heuristic.

Upper bound. We derive an upper bound of $V_{0,n}^{*}$ based on information structure relaxation. Consider a hypothetical model, where the information of demand type is revealed to the firm in the first period. In this case, the firm has *full information (fi)* about the demand type, and is able to decide the optimal capacity position contingent upon the demand type. Consequently, we obtain the firm's value-to-go function as follows:

$$V_{0,n}^{fi}(\boldsymbol{\pi_1}) = \max_{\mu_1,\dots,\mu_I} \sum_{i=1}^{I} \pi_{1,i} E\left\{ \sum_{j=l_n+1}^{J_n} \left[pD_j - c_1\left(D_j - n\mu_i\tau_n\right)^+ - c_0 n\mu_i\tau_n \right] - \hat{C}(n\mu_0, n\mu_i) \,\middle|\, \theta_{i,n} \right\} \quad (15)$$

We observe that the value-to-go function above is concave in the demand. Therefore, by Jensen's inequality, we have an upper bound of $V_{0,n}^{fi}$ from a *deterministic (d)* problem as follows:

$$\begin{aligned} V_{0,n}^{d}(\boldsymbol{\pi_1}) &= \max_{\mu_1,\dots,\mu_I} \sum_{i=1}^{I} \pi_{1,i} \left\{ \left[pn\lambda_i - c_1 n(\lambda_i - \mu_i)^+ - c_0 n\mu_i \right] (J_n - l_n)\tau_n - \hat{C}(n\mu_0, n\mu_i) \right\} \\ &= \sum_{i=1}^{I} \pi_{1,i} \left\{ (p - c_0) n\lambda_i (J_n - l_n)\tau_n - \hat{C}(n\mu_0, n\lambda_i) \right\} \end{aligned} \quad (16)$$

In the deterministic problem described in equation (16), the optimal target capacity for demand type i is $\mu_i^* = \lambda_i$. It is not a surprise that the firm's optimal action is to adjust the capacity to the mean instead of a newsvendor type fractile, because the decision problem is deterministic, and there is no uncertainty in the demand. Finally, it follows that $V_{0,n}^* \le V_{0,n}^{fi} \le V_{0,n}^d$.

Performance evaluation. To evaluate the performance of the policy in the asymptotic regime, we analyze the metric of *regret*, which measures the gap between the value-to-go function under the heuristic and the deterministic upper bound. Formally, the regret of the two-step heuristic is defined as $R_n^{ts} = 1 - V_{0,n}^{ts}/V_{0,n}^d$. In the following, we say a heuristic is asymptotically optimal if the regret converges to 0 as the scale factor n increases to infinity. For two sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n \asymp b_n$ if $a_n = O(b_n)$ and $b_n = O(a_n)$. Then we characterize the asymptotic regret as follows.

PROPOSITION 3 **(Asymptotic regret: Two-step heuristic)**.

If $\tau_n \asymp n^{-\frac{1}{3}}$ *for all* n, *the two-step heuristic is asymptotically optimal and* $R_n^{ts} = O\left(n^{-\frac{1}{3}}\right)$.

We first observe that the firm sets $\tau_n \asymp n^{-\frac{1}{3}}$ corresponding to a problem scale of n. This reflects the exploration-exploitation tradeoff the firm faces. For a given problem scale n, the firm has incentive to set a long observation period to *explore* the demand so that it can obtain more demand information. However, the longer the observation period is, the less time is left for the firm to *exploit* the benefit of its knowledge about demand by adjusting the capacity. Therefore, the firm has to choose an appropriate period length to balance this tradeoff. As the problem scale increases, more demand information is available within a unit of time. Therefore, the firm is able to reduce the observation period and starts to exploit its knowledge earlier. The result in Proposition 3 directly extends to the case where the variance does not depend on λ_i. We also extend the result to accommodate the case where the variance may increase in a different rate with respect to the problem scale n in Appendix EC.2.1.2.

5. Capacity Investment with Multiple Adjustment Opportunities

We now move on to the other case where the capacity adjustment cost does not depend on the number of adjustments the firm has committed to, i.e., the managerial hurdle parameter $\Lambda = 0$, and therefore the capacity adjustment cost $\hat{C}(\mu, \mu') \triangleq c_a(\mu' - \mu)^+ + \gamma_a(\mu - \mu')^+$. In this case, it follows that the firm has multiple opportunities to adjust capacity during a planning horizon. As described in Section 3, at the beginning of each period, the firm first decides whether it will adjust its capacity or not, and if so, by how much. Then the demand is realized and satisfied using the firm's capacity and (if short) an outside option. At the end of the period, the firm updates the posterior distribution of demand types. We first present the case where the capacity level set is discrete, i.e., $\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$, where $\mathbb{Z}_+$ stands for the set of nonnegative integers, and there are still I demand types. Then we also characterize a similar optimal policy for the continuous capacity level. We consider the case that the capacity adjustment decision is made based on the capacity position. We also discuss the case where capacity reduction cannot exceed the current installed capacity level in Section EC.2.3. We use a superscript m to indicate the multiple adjustment model.

Following equation (9), the optimal value-to-go function is recursively defined as follows: recall that $\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$, for all $j \in \{1, 2, ..., J - l\}$,

$$\begin{aligned} V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) &= \max_{\hat{\mu} \in \mathcal{K}} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - \hat{C}(\hat{\mu}_{j-1}, \hat{\mu}) + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu}) | \boldsymbol{\pi_j}\right] \right\}; \\ V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) &= 0 \text{ for } j > J - l. \end{aligned} \tag{17}$$

Let $\boldsymbol{\pi_j}$ be the information vector and $\hat{\mu}_{j-1}$ be capacity position at the beginning of period j. The next result shows that, for a given $\boldsymbol{\pi_j}$, the optimal policy is of a control band type.

PROPOSITION 4 **(Optimal policy for multiple adjustment opportunities)**. *The optimal capacity position, $\hat{\mu}^*(\boldsymbol{\pi_j})$, is characterized by two thresholds $\underline{\mu}_j(\boldsymbol{\pi_j})$ and $\bar{\mu}_j(\boldsymbol{\pi_j})$, such that:*
(i) If $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j})$, it is optimal to adjust the capacity position up to $\hat{\mu}^*(\boldsymbol{\pi_j}) = \underline{\mu}_j(\boldsymbol{\pi_j})$.
(ii) If $\underline{\mu}_j(\boldsymbol{\pi_j}) \le \hat{\mu}_{j-1} \le \bar{\mu}_j(\boldsymbol{\pi_j})$, it is optimal to stay put, i.e., $\hat{\mu}^*(\boldsymbol{\pi_j}) = \hat{\mu}_{j-1}$.
(iii) If $\hat{\mu}_{j-1} > \bar{\mu}_j(\boldsymbol{\pi_j})$, it is optimal to adjust the capacity position down to $\hat{\mu}^*(\boldsymbol{\pi_j}) = \bar{\mu}_j(\boldsymbol{\pi_j})$.

When proving the result for discrete capacity levels, the usual argument based on concavity does not work as the concavity requires the continuity. Instead, we use $L^\natural$-concavity (Murota 2003) and its Lovász extension, both of which have been used in inventory literature: see Zipkin (2008), Huh and Janakiraman (2010), Gong and Chao (2013) and Chen et al. (2014). To define $L^\natural$-concavity, let $\boldsymbol{e} \triangleq (1, ..., 1)$ be an n-dimensional vector of 1's and $S \subset \mathbb{Z}_+^n$ be a lattice. A function $f : S \to \mathbb{R}$ is $L^\natural$-concave if the function $\psi(\boldsymbol{v}, \xi) = f(\boldsymbol{v} - \xi\boldsymbol{e})$, a scalar $\xi \ge 0$, is supermodular on $\{(\boldsymbol{v}, \xi) | \boldsymbol{v} - \xi\boldsymbol{e} \in S\}$. In order to prove the result, we show that $L^\natural$-concavity is preserved in the dynamic program and its Lovász extension is a concave function. The proof is in Appendix EC.2.2.1.

This result expands the result of Eberly and Van Mieghem (1997) who show that the optimal policy is of a control band type for a given information vector to the discrete capacity case. Intuitively, as it is costly to adjust capacity, the firm will adjust the capacity only if the current capacity position is significantly lower or higher than the expected level. In contrast to the single adjustment case, where the value-to-go function is not necessarily concave in the initial capacity, we show that the value-to-go function is $L^{\natural}$-concave in the starting capacity for the multiple adjustment case. The $L^{\natural}$-concave value-to-go function and its Lovász extension enables a simple characterization of the optimal policy by two state-dependent capacity adjustment thresholds. Note that, however, the result only holds for a given information vector. Thus, in order to fully characterize the optimal policy, we also need to know how the optimal policy changes with respect to the information vector. Also, note that, in spite of the fact that the structure can be readily described, the optimal policy is still computationally complex because the information vector $\boldsymbol{\pi_j}$ is uncountable.

When the capacity type is continuous, i.e., $\mathcal{K} = \mathbb{R}_+$, the optimal value-to-go function is modified accordingly. The control band policy similar to Proposition 4 holds for the continuous capacity case. The details are relegated to the Appendix EC.2.2.1.

We next show how the two thresholds $\underline{\mu}_j(\boldsymbol{\pi_j})$ and $\bar{\mu}_j(\boldsymbol{\pi_j})$ change in the information vector $\boldsymbol{\pi_j}$. For this, we introduce the likelihood ratio order: for information vector $\boldsymbol{\pi_j} = (\pi_{j,1}, ..., \pi_{j,I})$ and $\boldsymbol{\pi'_j} = (\pi'_{j,1}, ..., \pi'_{j,I})$, we have $\boldsymbol{\pi_j}$ is smaller than $\boldsymbol{\pi'_j}$ in the likelihood ratio order sense, i.e., $\boldsymbol{\pi_j} \preceq_{lr} \boldsymbol{\pi'_j}$, if $\pi'_{j,i}/\pi_{j,i} \leq \pi'_{j,\hat{i}}/\pi_{j,\hat{i}}$ for any $i < \hat{i}$.

Proposition 5 **(Monotonicity of switching curves)**. *Both $\underline{\mu}_j(\boldsymbol{\pi_j})$ and $\bar{\mu}_j(\boldsymbol{\pi_j})$ increase as the information vector, $\boldsymbol{\pi_j}$, increases in the likelihood ratio order sense.*

Proposition 5 implies that as the underlying demand type becomes more likely to be high, both invest-up-to and divest-down-to thresholds increase. In other words, the switching curves and the resultant capacity levels that the firm sets under the optimal policy are both monotone in $\boldsymbol{\pi_j}$. This is a sharp contrast to the result of the single adjustment case, in which the optimal policy (and the resultant capacity level) is not monotone in the information vector. This highlights how optimal policy changes when the firm has very limited opportunities to change the capacity. In the single adjustment case, the firm needs to decide two things: *when to adjust* and *by how much.* As a result, the firm may decide to wait even in the state when $\boldsymbol{\pi_j}$ is high while it increases capacity when $\boldsymbol{\pi_j}$ is lower. In the multiple adjustment case, however, the firm does not need to worry about exhausting the opportunities (of course, whether to change or not, and to what level are still driven by capacity cost and leadtime). Although there is a stay-put interval, this is purely driven by cost. In the single adjustment case, however, the non-monotonicity is caused by the option value of adjustment opportunity (if we change capacity in this period, we cannot change it again).

Table 2 The multi-step heuristic

Given the period length τ_n and the number of adjustment opportunities K_n,

1. The firm serves the demand in period 1 with the initial capacity $n\mu_0$.
2. For $\kappa = 1 : K_n$
   a. The firm adjusts the capacity position at the start of period 2^κ to the observed average demand during the first $2^\kappa - 1$ periods denoted by $n\bar{\lambda}_\kappa$. The capacity level will be updated accordingly l_n periods later.
   b. The firm serves the demand from period 2^κ to $2^{\kappa+1} - 1$ using the (updated) capacity.
   End
3. The firm serves the demand in the remaining periods using the updated capacity.

Our results can be generalized to the case with fixed capacity cost, the capacity reduction is limited by the current available capacity, and different leadtimes to increase and decreases capacity. The details are provided in EC.2.2.2 and EC.2.3.

5.1. Near-Optimal Heuristic and Performance Evaluation

We next derive a simple near-optimal heuristic similar to the one in Section 4.2. The setting is entirely identical other than the fact that the firm is able to adjust its capacity multiple times during the decision horizon $[0, T]$ (equivalently period 1 to J_n in discrete time). We also show that in an asymptotic regime (same as the one defined in Section 4.2), this multi-step heuristic (ms) is asymptotically optimal, and provide a performance upper bound for the heuristic under the regret criterion. To show the asymptotic optimality, we also impose Assumption **??** on λ_i. Note that the choice of τ_n affects the performance of a heuristic. In our heuristic, we construct a sequence of τ_n to achieve the asymptotic optimality with a provable convergence rate. The result also holds for a demand process where the variance does not depend on λ_i.

The multi-step heuristic. In this heuristic, the firm adjusts its capacity only in a subset of the J_n periods, instead of doing it in every period. Specifically, the κ^{th} adjustment of the capacity position occurs at the beginning of period 2^κ, and the actual change of capacity levels occurs at the start of period $2^\kappa + l_n$, for $\kappa = 1, 2, ..., K_n$, where K_n is the largest integer such that $l_n + \sum_{\kappa=1}^{K_n+1} 2^{\kappa-1} \leq J_n$, i.e., $K_n \triangleq \lfloor \log_2(J_n - l_n + 1) \rfloor - 1$. That is, the time between the $\kappa - 1^{th}$ and κ^{th} adjustments is $2^{\kappa-1}\tau_n$ ($2^{\kappa-1}$ periods). The intuition for choosing the exponentially increasing periods between two consecutive adjustment decisions is that as more demand information is collected, adding new observations is less likely to change the information vector in a significant way. The details of the heuristic are illustrated in Table 2.

In this heuristic, the firm always adjusts the capacity position to the observed demand rate. To evaluate the value-to-go function under this heuristic, we denote the observed demand rate contingent upon the demand type i by $n\bar{\lambda}_{i,\kappa}$ for $\kappa \geq 1$. Then, we first define $\bar{\lambda}_{i,\kappa}$ recursively below.

$$\bar{\lambda}_{i,1} \triangleq \frac{D_1|\theta_{i,n}}{n\tau_n}$$

$$\bar{\lambda}_{i,\kappa} \triangleq \frac{\bar{\lambda}_{i,\kappa-1} n(2^{\kappa-1}-1)\tau_n + \sum_{j=2^{\kappa-1}}^{2^\kappa-1} D_j|\theta_{i,n}}{n(2^\kappa-1)\tau_n}, \ \kappa = 2,3,...,K_n \tag{18}$$

For notational simplicity, we also define $\bar{\lambda}_{i,0} \triangleq \mu_0$. Then we have the firm's expected value-to-go function under this heuristic as follows.

$$V_{0,n}^{ms}(\boldsymbol{\pi_1}) = \sum_{i=1}^{I} \pi_{1,i} E \left\{ \begin{array}{l} \sum_{\kappa=0}^{K_n-1} \sum_{j=l_n+2^\kappa}^{l_n+2^{\kappa+1}-1} \left[pD_j - c_1 \left(D_j - n\bar{\lambda}_{i,\kappa}\tau_n \right)^+ - c_0 n\bar{\lambda}_{i,\kappa}\tau_n \right] \\ + \sum_{j=l_n+2^{K_n}}^{J_n} \left[pD_j - c_1 \left(D_j - n\bar{\lambda}_{i,K_n}\tau_n \right)^+ - c_0 n\bar{\lambda}_{i,K_n}\tau_n \right] \\ - \sum_{\kappa=1}^{K_n} \hat{C}\left(n\bar{\lambda}_{i,\kappa-1}, n\bar{\lambda}_{i,\kappa} \right) \end{array} \middle| \theta_{i,n} \right\} \tag{19}$$

As this heuristic is a feasible policy for the corresponding optimal capacity adjustment problem, we have that $V_{0,n}^{ms} \le V_{0,n}^{m*}$, where $V_{0,n}^{m*}$ denotes the value-to-go function under the optimal policy. As the optimal policy is not computationally tractable, we need to derive an upper bound of the value-to-go function under the optimal policy in order to evaluate the performance of the heuristic.

Upper bound. We first observe that the $V_{0,n}^{d}$ (see equation (16) in Section 4.2) is still an upper bound of $V_{0,n}^{m*}$. This is because in the deterministic stationary demand setting, once the firm obtains full information about the demand type, even if the firm is able to adjust capacity any time, it is still optimal to adjust it only once at the beginning of the time horizon as the adjustment is costly. That is, we still have the optimal target capacity $\mu_i^* = \lambda_i$, and $V_{0,n}^{d}$ as follows.

$$V_{0,n}^{d} = \sum_{i=1}^{I} \pi_{1,i} \Big\{ (p-c_0)n\lambda_i(J_n - l_n)\tau_n - c_a n(\lambda_i - \mu_0)^+ - \gamma_a n(\mu_0 - \lambda_i)^+ \Big\} \tag{20}$$

Performance evaluation. To analyze the performance of the heuristic, we evaluate the asymptotic behavior of the regret of the multi-step heuristic, defined as $R_n^{ms} = 1 - V_{0,n}^{ms}/V_{0,n}^{d}$. We derive the following characterization of the asymptotic regret.

Proposition 6 **(Asymptotic regret: Multi-step heuristic)**.

If $\tau_n \asymp n^{-\frac{1}{3}}$ for all n, the multi-step heuristic is asymptotically optimal and $R_n^{ms} = O\left(n^{-\frac{1}{3}}\right)$.

The intuition of the proof is that as the firm observes more demand information and adjusts capacity to match the observed average demand rate, we are able to bound outsourcing costs and capacity adjustment costs by the bound shown in Proposition 1 in Gallego (1992), which derived a one-sided deviation bound for the class of distributions with finite mean and variance. As noted above, we choose the exponentially increasing time between two consecutive decisions because the adjustment is costly, and with more information learned, it is less necessary for the firm to learn about demand frequently. Finally, the time interval τ_n is set to minimize the derived upper bound.

Recall that when the firm has only one chance to adjust its capacity, the upper bound of the regret is also $O\left(n^{-1/3}\right)$ (see Proposition 3). Here, although the upper bound of the regret is still of the same order, the capacity adjustment cost makes a difference. With multiple adjustment opportunities, the firm is able to correct errors that it might have made in a one shot decision, and therefore, the regret should be smaller. However, when capacity adjustment is very costly, with multiple capacity adjustments specified in the heuristic, the firm needs to pay a higher total capacity adjustment cost as it chases the mean demand. Therefore, the benefit from the learning-while-doing may be diluted. In Section 6.2, we compare the two heuristics numerically.

6. Numerical Study

In order to demonstrate that our method can be implemented with actual demand data and a few parameters related to costs, we develop a numerical study where the model premises (such as demand pattern, problem scale, cost, and profit) are motivated from actual application. Although we need to make some simplifying assumptions because we estimated parameters from aggregate financial and accounting data, we show that the results and conclusions are quite robust to model parameters and our assumptions.

In what follows, we first illustrate the information we extracted from the production and financial data such as the demand pattern, and then explain the context and assumptions of our numerical study in section 6.1. In section 6.2, we use numerically examine the impact of the market size and leadtime to adjust capacity on the performance of our proposed heuristic. In addition, we compare the performance of the two-step heuristic versus the multi-step heuristic, and illustrate the performance of the heuristic relative to the optimal policy. We relegate the impact of misspecified demand and cost parameters to Appendix EC.3.

6.1. Data and Parameters

Our example utilizes production and financial data related to the Ford Focus and tries to solve the following problem. Using the data from the first two generations of the Focus in the North American market, the numerical study illustrates how one could use our heuristics in deciding how to adjust capacity for the third generation. In this section, we briefly describe how we collect data and estimate the demand and cost parameters, with details deferred to the appendix.

Demand. Our focus is on how the assembly factory should adjust its capacity based on orders received from the dealership. Therefore, the demand is closely related to the number of Ford Focus sedans produced at the assembly factory. While customers buy cars from dealers who hold inventory, the plant operation is close to a make-to-order environment as Ford produces cars based on dealer orders and its production plan.

To analyze the demand pattern, we first collect monthly production data of Ford Focus in North America from January 2005 to December 2010 from the database of *Automotive News Data Center*[4]. There are two (redesigned) generations of Focus during this period: the first from January

Figure 3 **Monthly production of Ford Focus in North America market (a), and the empirical cdf and gamma distribution with sample mean and variance (b), from January 2005 to December 2010**

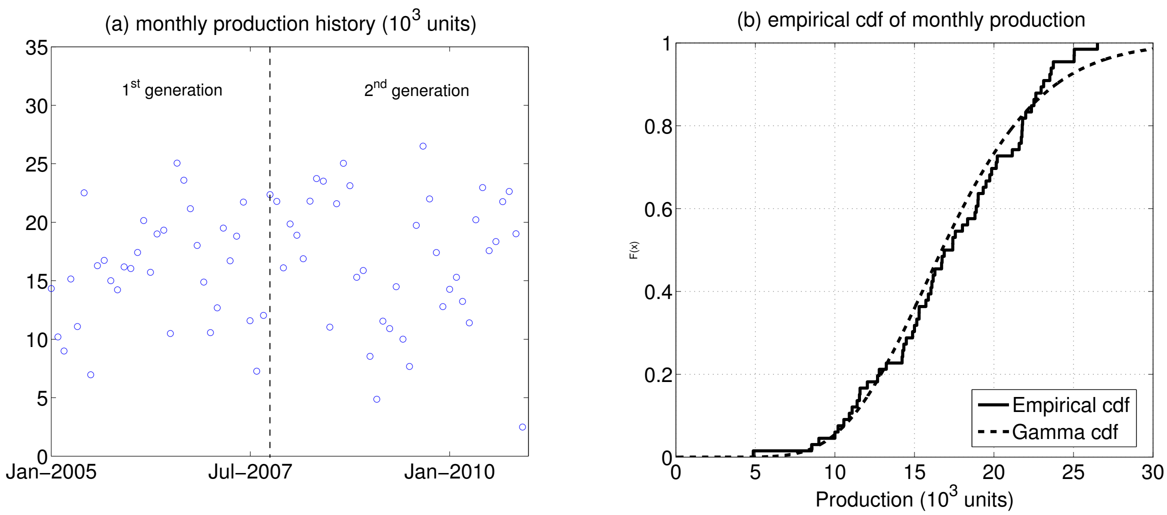

2005 to September 2007, and the second from October 2007 to December 2010. Although there is seasonality within each year affected by factors such as mid-year discount when manufacturers switch production to the next year model and end-of-year sales to boost sales figure, we observe the demand pattern is plausibly *stationary* within each generation: see Figure 3(a). As the demand pattern of the first two generations is similar, we group the production data from January 2005 to December 2010 and observe that the monthly demand approximately follows a gamma distribution with a mean of 17.21 thousand units per month and a standard deviation of 5.04; see Figure 3(b).[5] When we construct the empirical cumulative distribution function, we excluded the data point in July or August to account for the regular summer shutdown. We also excluded the data when there was a model transition in that month. We denote the two key parameters of the gamma distribution by a and b, i.e., the probability density function is characterized as

$$f(x|a,b) = \frac{1}{\Gamma(a)b^a} x^{a-1} e^{-\frac{x}{b}}$$

where $\Gamma(\cdot)$ represents the gamma function.

In fact, we test the cleaned production data from January 2005 to December 2010 with a gamma distribution where the estimated parameters are $a = 11.67$ and $b = 1.47$ using a one-sample Kolmogorov-Smirnov test, which yields a p-value of 0.74, supporting our choice for the demand distribution. Therefore, we model the monthly demand (with the unit of a thousand cars) for Focus using a stationary gamma distribution.

Based on the demand pattern extracted from the production data, we next construct the context for the numerical study. We would like to note that our heuristic, which is data driven, does not rely on knowledge about the prior distribution or the exact demand distribution for each type.

The specific context of our numerical study is necessary only to evaluate the performance of the heuristic, i.e., computing the regret. Also, it should be noted that our heuristic applies to more general settings, e.g., when there are more than three market scenarios, or the unknowns are a vector of parameters rather than a single parameter.

To evaluate the performance of our heuristic, we postulate the following scenario. The decision maker has three possible scenarios (demand types) for the third generation of the product, low, medium and high. In the medium scenario, the demand will remain at the same level as the first two generations: monthly demand will follow the same distribution. In the other two scenarios, the demand for the third generation (released in May 2011) is either lower or higher than the first two generations depending on the popularity of the third generation and economic environment.

In the low and high scenarios, we assume the average monthly demand is either dropped by or raised by 5 thousand units (which is about one standard deviation). That is, the two key parameters are $a = 8.28$ and $b = 1.47$ for the low demand case, and $a = 15.06$ and $b = 1.47$ for the high demand case. We assume that the parameter b, which stands for the ratio between variance and mean, stays stationary, i.e., a higher demand is associated with a higher variance. We will later show that our result is quite robust with respect to the misspecification of the average demand parameters.

Finally, as the second generation Focus was on sale for three years, we assume the decision horizon T for the third generation is also 3 years, starting from January 2011. Following the convention of the asymptotic analysis, we also assume when $n = 1$, the average medium type demand in the three year horizon is 1 unit and $\tau_1 = 36$ months. Therefore, the problem scale in the base case is $n_1 = 17.21 \times 10^3 \times 36 = 619,610$, and we assume the firm reviews demand and makes capacity adjustment decisions in a monthly scale at the current demand level, i.e., $\tau_{n_1} = 1$ (recall that $\tau_{n_1} \asymp n_1^{-1/3}$ and the practical limitation that the data used to construct the numerical study were collected on a monthly scale.) We will illustrate the impact of market size on the performance of the heuristic in the numerical study. We also assume there is no leadtime, i.e., $l = 0$, and we will study the impact of leadtime later.

Initial capacity. Our target is to analyze Ford's capacity adjustment decision for the third generation. Therefore, besides demand information, we also need capacity information. Since Ford does not publish their exact maximum capacity, we use the maximum production quantity from January 2010 to December 2010 as the starting capacity, i.e., 22.97 thousand cars per month.[6]

Cost/profit parameters. We use aggregated cost parameters at the firm level to approximate the ones at the product level. Specifically, we recover the gross capacity of Ford using its public financial reports and data, and then identify the unit profit and capacity related costs at the firm level. Although these are rough estimates, the performance of our heuristic is quite robust to the cost parameters. Note that the cost parameters would be significantly more accurate if one could

Table 3 Production and capacity related profit/cost parameters

Estimated cost parameters for Focus	Value
Capacity adjustment (upgrading) cost c_a	\$4,487 *month/unit*
Capacity adjustment (downgrading) cost γ_a	\$448.7 *month/unit*
Capacity overhead cost c_0	\$181.1 *per unit*
Capacity outsourcing cost c_1	\$362.2 *per unit*
Unit profit p (excluding capacity related cost)	\$1,270 *per unit*
Average retail price	\$22,154 *per unit*

extract the cost information from an ERP or internal accounting system. We summarize the cost parameters derived from Ford's Annual Report in 2012 (Ford Motor Company 2012) in Table 3, and relegate the details of estimations to Appendix EC.3.

6.2. Numerical Analysis

In the numerical study, we evaluate the performance of our heuristics using the regret with respect to its deterministic upper bound. Specifically, for given scale parameter n, define $R_n^{ts} = 1 - V_{0,n}^{ts}/V_{0,n}^{d}$ and $R_n^{ms} = 1 - V_{0,n}^{ms}/V_{0,n}^{d}$ to be regrets associated with the two-step heuristic and multi-step heuristic, respectively. The decision horizon, demand distributions, initial capacity, and profit and cost parameters are the ones specified in Section 6.1. In what follows, we first present the impact of various parameters and demand assumptions (market size and leadtime) on the performance of our two-step heuristic. We then compare the performance of the two-step heuristic with the multi-step heuristic. Finally, we show the performance of the two-step heuristic with respect to the optimal policy. We relegate the impact of misspecified demand, cost and profit parameters to Appendix EC.3 for the interest of space.

To evaluate the value-to-go function under the two-step heuristic, $V_{0,n}^{ts}$, for a given prior vector $\boldsymbol{\pi_1}$, we apply a simulation approach with 10^6 experiments. In each round, a demand distribution (a demand type) is first generated according to the prior, then a sample path of demand in each period is generated according to the distribution. For each sample path generated, the firm follows the two-step heuristic, and the resultant profit is calculated. We use the average of the 10^6 observations to approximate $V_{0,n}^{ts}$. The deterministic upper bound, $V_{0,n}^{d}$, is computed following equation (16).

Market size. We first analyze the impact of market size, which is determined by the scale factor n. From Proposition 3, when the scale factor of the decision problem is n, setting the length of the learning period as $\tau_n \asymp n^{-1/3}$ results in asymptotic convergence at most on the order of $n^{-1/3}$. In Figure 4, we show that as $log(n)$ increases linearly, the log of the regret decreases linearly. In the base case ($n_1 = 619,610$), we assume the firm reviews the demand information monthly, and adjusts capacity based on the observation in the first month, i.e., $\tau_{n_1} = 1$ month, following earlier discussions in Section 6.1. To analyze the impact of the market size, we let n be

Figure 4 **Regret with respect to market size when the prior $\boldsymbol{\pi_1} = (0.2, 0.4, 0.4)$**

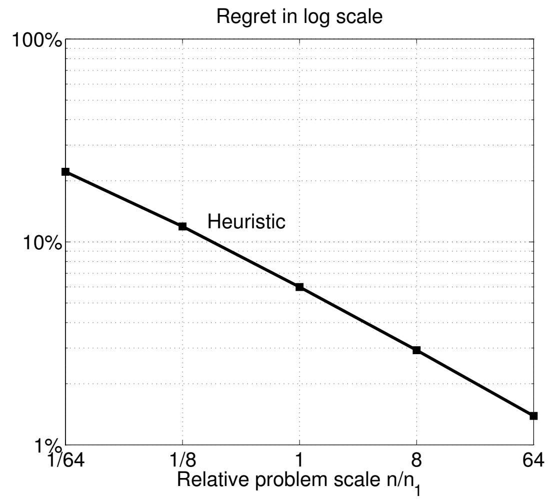

Figure 5 **Regret and value-to-go function with respect to leadtime when the prior $\boldsymbol{\pi_1} = (0.2, 0.4, 0.4)$**

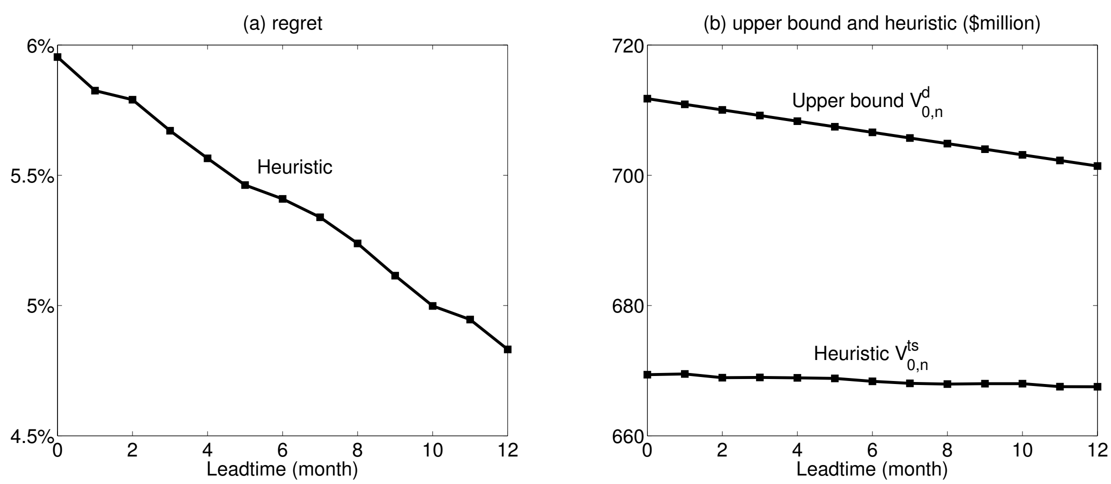

$8^k n_1, k = -2, -1, ..., 2$, corresponding to a τ_n of $2^{-k}, k = -2, -1, ..., 2$ month respectively. That is, as the magnitude of demand increases, the firm can adjust capacity within a smaller window of demand data. For instance, when $k = -2$, τ_n is 4 months, and when $k = 2$, τ_n is $1/4$ months (about 1 week). In Figure 4, we observe that the log of regret decreases at the slope of -0.33, corresponding to the $n^{-1/3}$ convergence rate. This implies that the absolute difference between the upper bound and the heuristic is sub-linear in n. The cases are similar when the priors are different, so for the interest of space the details are not shown here.

Leadtime. In the base case, we normalize the leadtime as 0. One may think that this might favor the two-step heuristic, but the result is the opposite. The performance of the two-step heuristic improves as the leadtime becomes longer. To show this, we change the leadtime l from 0 to 12 months when the review period τ_{n_1} is 1 month and compute the total revenue of the planning

Figure 6 **Two-step heuristic vs. multi-step heuristic when the prior $\boldsymbol{\pi_1} = (0.2, 0.4, 0.4)$**

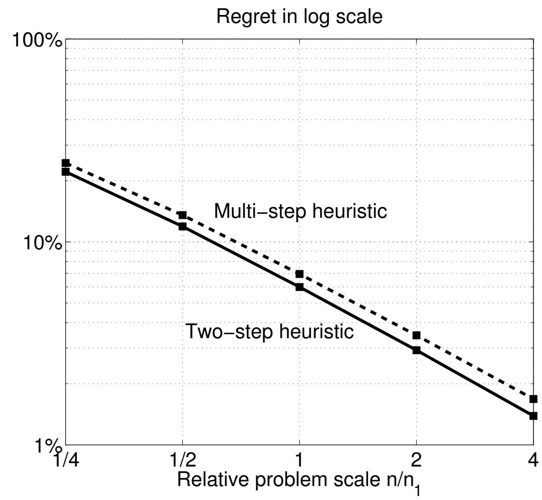

horizon (i.e., the value function plus the revenue of the first l periods (before any adjustment is made). In Figure 5(a), we observe that the regret decreases as the leadtime increases: the relative profit loss due to the lack of information decreases in leadtime. Although this is counter-intuitive at first glance, we observe from Figure 5(b) that, with a longer leadtime, the benefit of full information decreases, thus the performance of the deterministic upper bound deteriorates substantially, resulting in the decrease in the regret.

Single vs. multiple adjustments. We now compare the two-step heuristic with the multi-step heuristic. The capacity adjustment cost is specified in Section 6.1. In Figure 6, we observe that as the market size increases, the regrets of both policies decrease. In this case, as the firm needs to pay a much higher adjustment cost under the multi-step heuristic, which dominates the benefit from extra opportunities to adjust capacity, we observe that the regret under the multi-step heuristic is higher than the one under the two-step heuristic. However, when the capacity adjustment cost is small, as one may expect, the regret under the multi-step heuristic is lower than the one under the two-step heuristic, which reflects the benefit of learning-while-doing.

Heuristic vs. optimal policy. To simplify the computation for the optimal policy, we consider only two demand types in this part: *medium* and *high*. As there are only two demand types, we use π_j, the posterior distribution of high demand, to denote the information vector. Figure 7 shows the regret of the two-step heuristic. Compared to the deterministic upper-bound (which assumes the knowledge of full information and no randomness), the regret of our data-driven heuristic is no more than 6.03%. We use the deterministic upper bound to define the regret, because a large state-space makes it intractable to compute the optimal policy and resultant value function. In the two demand-type case, however, we can numerically approximate the value function of the optimal policy, $V^*_{0,n}$, with linear interpolation (i.e., evaluating the value at a set of fine fixed grid points

Figure 7 **Regret of the two-step heuristic with respect to upper bound and optimal policy**

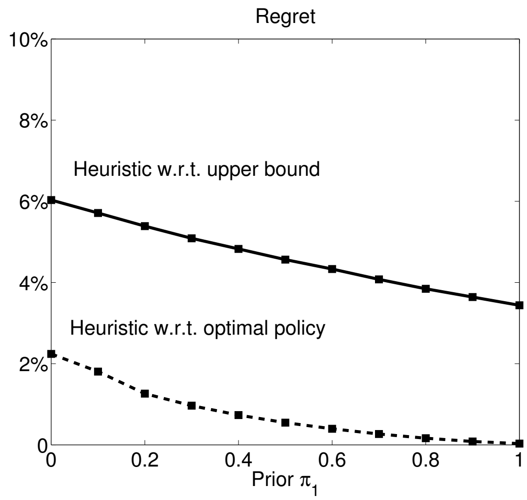

Note. There are only two demand types here and π_1 indicates the prior of the demand being high type

Figure 8 **Firm's capacity decision under the two-step heuristic and optimal policy**

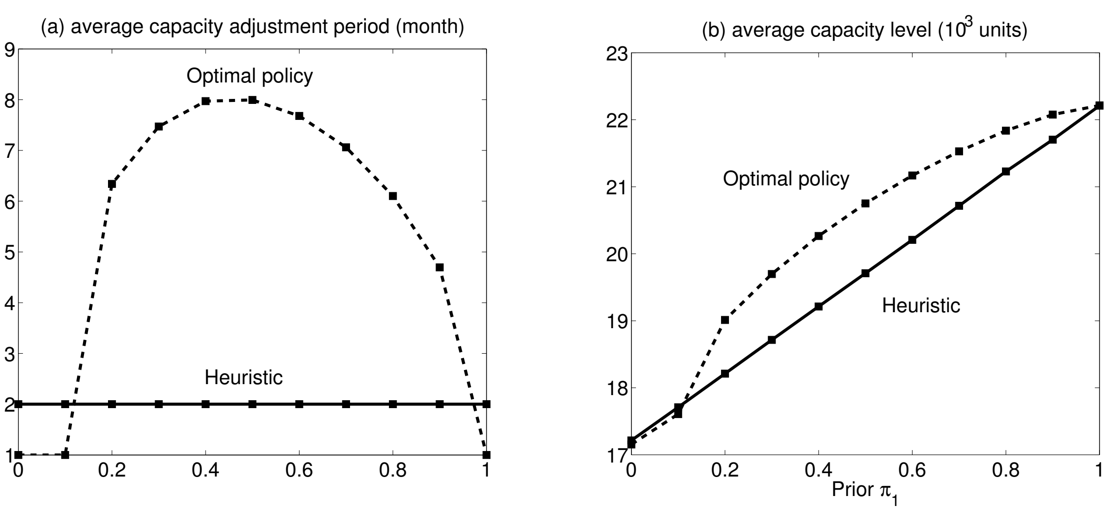

and then approximating values for the rest of the states using linear interpolation). As Figure 7 shows, the regret (compared to the optimal policy) is less than 2.24%.

From the timing perspective (Figure 8(a)), the firm always adjusts its capacity in the second month under the two-step heuristic. On the other hand, under the optimal policy the firm adjusts capacity early (in the first period) when the prior is close to the extremes (π_1 close to 0 or 1), and delays the decision when there is no dominant demand type in the prior. In addition, Figure 8(b) shows that, on average, the capacity levels under the optimal policy and the two-step heuristic are fairly close when the firm adjusts the capacity at the beginning of the decision period, because the optimal capacity level in this case is close to the average demand. When the firm is less certain about the demand type and prefers to delay the capacity adjustment to the future, consistent with the conventional wisdom, the firm invests more in capacity compared to the average capacity level built under the two-step heuristic.

7. Conclusions

We analyze a firm's capacity investment decision for a product with a finite planning horizon, and investigate ***when***, and by ***how much***, the firm should adjust its capacity. When the capacity adjustment costs increase significantly with respect to the number of adjustments and therefore the firm can adjust the capacity once in a planning horizon, we show that the firm may alternate its decision to pull the trigger (adjust capacity) or delay the adjustment multiple times as the likelihood of demand type changes. Although the optimal policy in general is non-monotone in the likelihood, we characterize the underlying structure of optimal policy. We show that if the firm decides to adjust the capacity, the target capacity position increases in the likelihood. We demonstrate that, even after knowing the structure of optimal policy, computing and implementing the policy can be very difficult. Instead, we present a very simple but provably well-performing data-driven heuristic when demand follows a stochastic process with stationary and independent increment. In this heuristic, the firm observes demand during an exploration period, and then adjusts capacity to match the observed demand rate. By carefully choosing the length of exploration period, the firm is able to balance the exploration and exploitation tradeoff, and the regret of the heuristic asymptotically converges to 0.

When the capacity adjustment costs remain stationary with respect to the number of adjustments and therefore the firm has multiple opportunities to adjust capacity, we show the firm's optimal policy is a control band policy, characterized by thresholds. Under this policy, in each period, the firm stays put to observe the demand when the capacity is between the two thresholds, and adjusts its capacity to the lower threshold only when the capacity is below it, and vice versa. We also develop a simple but asymptotically optimal heuristic, in which the firm predetermines a set of time points at which the firm will adjust its capacity to match the observed demand rate. The time between two consecutive decisions increases exponentially, reflecting the fact that the adjustment is costly, and it is less necessary for the firm to adjust capacity frequently with more demand information collected. The multiple adjustments enable the firm to correct errors in early decisions. However, when the capacity adjustment cost is high, the multiple adjustments also yield a higher adjustment cost, which dilutes the benefit of the learning-while-doing. The optimal policy and heuristics are illustrated using a numerical study.

To the best of our knowledge, this is one of the first papers analyzing the capacity adjustment with demand learning by characterizing the optimal policy as well as deriving near-optimal heuristics. We believe there are more opportunities in the area combing learning and capacity management. For example, how does the firm's learning opportunity affect the joint decision of capacity and inventory? How should the firm compute its capacity management strategy efficiently when the product life cycle is not stationary? We leave these questions for future research opportunities.

Endnotes

1. Since each period is τ units of time, the maximum demand the firm can satisfy with its own capacity in this case is $\mu\tau$.

2. The detailed derivation of the value-to-go functions is provided in Section EC.1.

3. We note that our results do not require $\{N(t), t \geq 0\}$ to be weakly increasing. In particular, when $\{N(t), t \geq 0\}$ represents a Brownian motion with positive drift, it implies that the demand in each period forms a sequence of i.i.d. normal random variables. Although this implies that the realized demand may be negative, assumptions of demands following normal distribution have been used for its analytical tractability, c.f., Eppen (1979) and Dong and Rudi (2004), particularly when the standard deviation is small relative to the mean.

4. Automotive News Data Center: `http://www.autonews.com/section/datacenter`.

5. We observe that during the automotive industry crisis (2008-2010), the demand pattern of Focus did not change. This may be because the Focus is a fuel-efficient model, and therefore the substantial increase in the prices of automotive fuels did not cause a significant drop in sales, unlike the sport utility vehicles and pickup trucks, whose demands declined in the same period.

6. According to Ford Motor Company (2012), the vehicle assembly capacity is categorized as *installed capacity* and *manned capacity*. Installed capacity refers to "the physical capability of a plant and equipment to assemble vehicles if fully manned". Manned capacity refers to "the degree to which the installed capacity has been staffed". In this numerical example, we use capacity to refer to the *installed* capacity that is specific to Ford Focus.

References

Aviv, Y., A. Pazgal. 2005. A partially observed markov decision process for dynamic pricing. *Management Sci.* **51**(9) 1400–1416.

Azoury, K. S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* **31**(9) 1150–1160.

Besanko, D., U. Doraszelski, L. X. Lu, M. Satterthwaite. 2010. Lumpy capacity investment and disinvestment dynamics. *Oper. Res.* **58**(4-Part-2) 1178–1193.

Besbes, O., A. Muharremoglu. 2013. On implications of demand censoring in the newsvendor problem. *Management Sci.* **59**(6) 1407–1424.

Besbes, O., A. Zeevi. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.* **57**(6) 1407–1420.

Boyacı, T., Ö. Özer. 2010. Information acquisition for capacity planning via pricing and advance selling: When to stop and act? *Oper. Res.* **58**(5) 1328–1349.

Burnetas, A., S. Gilbert. 2001. Future capacity procurements under unknown demand and increasing costs. *Management Sci.* **47**(7) 979–992.

Burnetas, A. N., C. E. Smith. 2000. Adaptive ordering and pricing for perishable products. *Oper. Res.* **48**(3) 436–443.

Chao, X., H. Chen, S. Zheng. 2009. Dynamic capacity expansion for a service firm with capacity deterioration and supply uncertainty. *Oper. Res.* **57**(1) 82–93.

Chen, L., E. L. Plambeck. 2008. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Oper. Management* **10**(2) 236–256.

Chen, W., M. Dawande, G. Janakiraman. 2014. Fixed-dimensional stochastic dynamic programs: An approximation scheme and an inventory application. *Oper. Res.* **62**(1) 81–103.

Davis, M. H. A., M. A. H. Dempster, S. P. Sethi, D. Vermes. 1987. Optimal capacity expansion under uncertainty. *Advances in Applied Probability* **19**(1) 156–176.

Dixit, A. K., R. S. Pindyck. 1994. *Investment under uncertainty*. Princeton University Press, Princeton, NJ.

Dong, L., N. Rudi. 2004. Who benefits from transshipment? exogenous vs. endogenous wholesale prices. *Management Sci.* **50**(5) 645–657.

Eberly, J. C., J. A. Van Mieghem. 1997. Multi-factor dynamic investment under uncertainty. *J. Econom. Theory* **75**(2) 345–387.

Power Assure. 2009. Dynamic power management: Adjusting data center capacity in real-time. *2009 Data Center Efficiency Summit Case Studies, Silicon Valley Leadership Group* `http://svlg.org/wp-content/uploads/2012/12/PowerAssure\_cs.pdf`. Retrieved May 27, 2015.

Eppen, G. D. 1979. Note-effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* **25**(5) 498–501.

Eppen, G. D., A. V. Iyer. 1997. Improved fashion buying with bayesian updates. *Oper. Res.* **45**(6) 805–819.

Ford Motor Company. 2012. Profitable growth for all: Ford Motor Company 2012 annual report `http://corporate.ford.com/doc/ar2012-2012%20Annual%20Report.pdf`. Retrieved June 1, 2013.

Freidenfelds, J. 1981. *Capacity expansion: analysis of simple models with applications*. Elsevier North Holland, New York.

Gallego, G. 1992. A minmax distribution free procedure for the (Q,R) inventory model. *Oper. Res. Lett.* **11**(1) 55–60.

Gong, X., X. Chao. 2013. Technical note-optimal control policy for capacitated inventory systems with remanufacturing. *Oper. Res.* **61**(3) 603–611.

Huh, W. T., G. Janakiraman. 2010. On the optimal policy structure in serial inventory systems with lost sales. *Oper. Res.* **58**(2) 486–491.

Huh, W. T., R. Levi, P. Rusmevichientong, J. B. Orlin. 2011. Adaptive data-driven inventory control with censored demand based on Kaplan-Meier estimator. *Oper. Res.* **59**(4) 929–941.

Huh, W. T., P. Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Math. Oper. Res.* **34**(1) 103–123.

Kaminsky, P., M. Yuen. 2014. Production capacity investment with data updates. *IIE Transactions* **46**(7) 664–682.

Kwon, H. D., S. A. Lippman. 2011. Acquisition of project-specific assets with Bayesian updating. *Oper. Res.* **59**(5) 1119–1130.

Lariviere, M. A., E. L. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* **45**(3) 346–363.

Lovejoy, W. S. 1991. A survey of algorithmic methods for partially observed markov decision processes. *Ann. Oper. Res.* **28**(1) 47–65.

Lovejoy, W. S. 1993. Suboptimal policies, with bounds, for parameter adaptive decision processes. *Oper. Res.* **41**(3) 583–599.

Luss, H. 1982. Operations research and capacity expansion problems: A survey. *Oper. Res.* **30**(5) 907–947.

Manne, A. S. 1967. *Investments for capacity expansion, size, location, and time-phasing*. MIT Press, Cambridge, MA.

Monahan, G. E. 1982. State of the arta survey of partially observable markov decision processes: Theory, models, and algorithms. *Management Sci.* **28**(1) 1–16.

Murota, K. 2003. *Discrete convex analysis*. SIAM.

Scarf, H. 1959. Bayes solutions of the statistical inventory problem. *Ann. Math. Statist.* **30**(2) 490–508.

Snow, D. C., S. C. Wheelwright, A. B. Wagonfield. 2006. Genentech–capacity planning. Case Study, Harvard Business Publishing.

Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Oper. Management* **5**(4) 269–302.

Wang, W., M. Ferguson, S. Hu, G. C. Souza. 2013. Dynamic capacity investment with two competing technologies. *Manufacturing & Service Oper. Management* Forthcoming.

Williams, D. 1991. *Probability with martingales*. Cambridge university press.

Zipkin, P. 2008. On the structure of lost-sales inventory models. *Oper. Res.* **56**(4) 937–944.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Supplement Materials for "Capacity Investment with Demand Learning"

EC.1. Derivation of the value-to-go functions in the single adjustment case

In this section, we provide an alternative detailed derivation of the firm's value-to-go function in Section 4. To model the firm's capacity decision, we first introduce the state vector $\boldsymbol{\omega_j} = (\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1})$. Here, $\boldsymbol{\pi_j}$ is the firm's belief about demand type given the demands up to period $j-1$, and $\hat{\mu}_{j-1}$ is defined as the capacity position in period $j-1$ (the capacity position $\hat{\mu}_{j-1}$ represents the capacity level in period $j+l-1$ since the capacity leadtime is l periods. In general, for period k, we have $\hat{\mu}_k = \mu_{k+l}$ and $\mu_k = \hat{\mu}_{k-l}$). Lastly, v_{j-1} is defined as an indicator to denote whether capacity has been changed on or prior to period $j-1$. Formally, if capacity adjustment is made in period j, we define

$$v_k = \begin{cases} 0 & \text{if } k < j \\ 1 & \text{if } k \geq j \end{cases} \tag{EC.1}$$

We next describe the transition of the state vector. We first observe that the transition of $\boldsymbol{\pi_j}$ is specified in equation (1) and $\boldsymbol{\pi_j}$ satisfies the property in Lemma 1. To describe how capacity position changes, we first introduce u_j to represent the firm's decision to adjust capacity in period j:

$$u_j = \begin{cases} 0 & \text{if the firm decides to stay put and continue to observe the demand} \\ 1 & \text{if the firm decides to adjust capacity in period } j \end{cases} \tag{EC.2}$$

As the firm has only a single opportunity to adjust the capacity, the feasible action space to adjust capacity in period j for given v_{j-1}, $\mathcal{A}(v_{j-1})$, is contingent upon whether the firm has adjusted the capacity or not, i.e.,

$$\mathcal{A}(v_{j-1}) = \begin{cases} \{0,1\} & \text{if } v_{j-1} = 0; \\ \{0\} & \text{if } v_{j-1} = 1. \end{cases} \tag{EC.3}$$

If $u_j = 1$, the firm adjusts the capacity level from the initial level μ_0 to maximize the expected profit from period $j+l$ till the end of the planning horizon based on the information vector $\boldsymbol{\pi_j}$. The induced target capacity position $\hat{\mu}_j^a(\boldsymbol{\pi_j})$, i.e., the capacity level that maximizes the remaining profit from period $j+l$ till the end is specified as follows.

$$\hat{\mu}_j^a(\boldsymbol{\pi_j}) \triangleq \arg\max_{\mu \in \mathcal{K}} E\left[\sum_{k=j+l}^{J} h_k(\boldsymbol{\Pi_k}, \mu) - \hat{C}(\mu_0, \mu) \middle| \boldsymbol{\pi_j} \right] = \arg\max_{\mu \in \mathcal{K}} \left\{ \sum_{k=j+l}^{J} h_k(\boldsymbol{\pi_j}, \mu) - \hat{C}(\mu_0, \mu) \right\}. \tag{EC.4}$$

The equality follows Lemma 1 and the fact that $h_k(\boldsymbol{\Pi_k}, \mu)$ is linear in $\boldsymbol{\Pi_k}$. When the maximizer is not unique, as a tie-breaking rule, the firm chooses the smallest capacity level. Then, the firm's (induced) capacity position transits as follows.

$$\hat{\mu}_j(\boldsymbol{\omega_j}, u_j) = \begin{cases} \hat{\mu}_j^a(\boldsymbol{\pi_j}) & \text{if } u_j = 1; \\ \hat{\mu}_{j-1} & \text{if } u_j = 0. \end{cases} \tag{EC.5}$$

We observe that using $\hat{\mu}_j^a(\boldsymbol{\pi_j})$ defined in equation (EC.4) in the dynamic program turns the firm's decision problem into an optimal stopping problem, i.e., when to pull the trigger and adjust the capacity to the level specified by $\hat{\mu}_j^a(\boldsymbol{\pi_j})$. For ease of exposition, we suppress the dependence of $\hat{\mu}_j^a(\boldsymbol{\pi_j})$ on $\boldsymbol{\pi_j}$ when there is no confusion.

Having characterized the state transition, we next define the objective function. Given the indicator v_{j-1}, and the starting capacity position $\hat{\mu}_{j-1}$, if the firm adjusts its capacity position to $\hat{\mu}_j^a$ in period j (i.e., $u_j=1$), with a capacity leadtime of l periods, it accrues profit in period $j+l$ with capacity $\mu_{j+l}=\hat{\mu}_j^a$, but pays a capacity adjustment cost in period j. Otherwise the firm's capacity level in period $j+l$ will be $\hat{\mu}_{j-1}$. Formally, we have the expected operating profit in period $j+l$ minus any capacity cost that the firm incurred, $H_j(\boldsymbol{\pi_j},\hat{\mu}_{j-1},v_{j-1},u_j)$, as

$$\begin{aligned}
H_j(\boldsymbol{\pi_j},\hat{\mu}_{j-1},v_{j-1},u_j) &\triangleq E[h_{j+l}(\boldsymbol{\Pi_{j+l}},\hat{\mu}_j(\boldsymbol{\omega_j},u_j)) - \hat{C}\left(\hat{\mu}_{j-1},\hat{\mu}_j(\boldsymbol{\omega_j},u_j)\right)|\boldsymbol{\pi_j}] \\
&= h_{j+l}(\boldsymbol{\pi_{j+l}},\hat{\mu}_j(\boldsymbol{\omega_j},u_j)) - \hat{C}\left(\hat{\mu}_{j-1},\hat{\mu}_j(\boldsymbol{\omega_j},u_j)\right) \\
&= \begin{cases} h_{j+l}(\boldsymbol{\pi_j},\hat{\mu}_j^a) - \hat{C}\left(\mu_0,\hat{\mu}_j^a\right) & \text{if } u_j=1 \\ h_{j+l}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) & \text{if } u_j=0 \end{cases}
\end{aligned} \tag{EC.6}$$

The first equality follows Lemma 1 and the fact that $h_{j+l}(\boldsymbol{\Pi_{j+l}},\hat{\mu}_j)$ is linear in $\boldsymbol{\Pi_{j+l}}$. For ease of exposition, we suppress the dependency of $\hat{\mu}_j(\boldsymbol{\omega_j},u_j)$ on $\boldsymbol{\omega_j}$ and u_j when there is no confusion.

To represent the firm's capacity decision as a dynamic program, we define a *policy* as a sequence of functions mapping the state vector to the action space $\mathcal{A}(v_{j-1})$ for all $j\le J-l$, i.e., $\{u_j(\boldsymbol{\omega_j}), j=1,2,...,J-l\}$. We notice that with a leadtime of l, the firm should not adjust its capacity after period $J-l$. Let $\mathcal{G}$ denote the set of all the admissible policies, and the firm's objective is to find a policy $g^*\in\mathcal{G}$ to maximize the expected total profit,

$$\max_{g\in\mathcal{G}}\sum_{k=1}^{l}E\left[h_k(\boldsymbol{\Pi_k},\mu_0)|\,\boldsymbol{\pi_1}\right]+\sum_{k=1}^{J-l}E^g\left[H_k(\boldsymbol{\Pi_k},\hat{\mu}_{k-1},v_{k-1},u_k)|\,\boldsymbol{\pi_1}\right] \tag{EC.7}$$

where the expectation is taken over D_j for all j at time zero. Due to the l-period leadtime, the expected profit of the first l periods, $\sum_{k=1}^{l}E\left[h_k(\boldsymbol{\Pi_k},\mu_0)|\,\boldsymbol{\pi_1}\right]$, is independent of the firm's capacity adjustment policy. Therefore, it is sufficient to maximize

$$\max_{g\in\mathcal{G}}\sum_{k=1}^{J-l}E^g\left[H_k(\boldsymbol{\Pi_k},\hat{\mu}_{k-1},v_{k-1},u_k)|\boldsymbol{\pi_1}\right] \tag{EC.8}$$

Define a partial policy $g_j\triangleq\{u_k(\boldsymbol{\pi_k},\hat{\mu}_{k-1},v_{k-1}),k=j,...,J-l\}$ and the set of all the admissible partial policies by $\mathcal{G}_j$. Then at the beginning of period j, given the initial states $\boldsymbol{\pi_j}$, $\hat{\mu}_{j-1}$ and v_{j-1}, the firm's optimal value-to-go function is

$$V_j(\boldsymbol{\pi_j},\hat{\mu}_{j-1},v_{j-1})=\max_{g_j\in\mathcal{G}_j}\sum_{k=j}^{J-l}E^{g_j}\left[H_k(\boldsymbol{\Pi}_k,\hat{\mu}_{k-1},v_{k-1},u_k)|\,\boldsymbol{\pi_j}\right] \tag{EC.9}$$

Then, the optimal value-to-go functions satisfy the following recursive optimality equations for all $j \in \{1, 2, ..., J-l\}$.

$$V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1}) = \max_{u_j \in \mathcal{A}(v_{j-1})} \{H_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1}, u_j) + E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_j, v_j) | \boldsymbol{\pi_j}\right]\}$$
$$V_k(\boldsymbol{\pi_k}, \hat{\mu}_k, v_k) = 0, \text{ for } k > J-l \tag{EC.10}$$

To simplify the optimality equations above, we observe the following: for $j = 1, 2, ..., J-l$, if the firm has not adjusted the capacity before period j, i.e., $v_{j-1} = 0$, we have $\hat{\mu}_{j-1} = \mu_0$. In this case, if the firm decides to adjust its capacity in period j, i.e., $u_j = 1$, then for $k = j+1, ..., J-l$, we have $\mathcal{A}(v_{k-1}) = \{0\}$ and $u_k = 0$, and therefore the firm's value-to-go function is as follows:

$$\begin{aligned} L_j^a(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1}) = L_j^a(\boldsymbol{\pi_j}, \mu_0, 0) &\triangleq h_{j+l}\left(\boldsymbol{\pi_j}, \hat{\mu}_j^a\right) - \hat{C}\left(\mu_0, \hat{\mu}_j^a\right) + E\left[V_{j+1}\left(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_j^a, 1\right) \middle| \boldsymbol{\pi_j}\right] \\ &= h_{j+l}\left(\boldsymbol{\pi_j}, \hat{\mu}_j^a\right) - \hat{C}\left(\mu_0, \hat{\mu}_j^a\right) + \sum_{k=j+1}^{J-l} H_k\left(\boldsymbol{\pi}_j, \hat{\mu}_j^a, 1, 0\right) \\ &= \sum_{k=j+l}^{J} h_k\left(\boldsymbol{\pi_j}, \hat{\mu}_j^a\right) - \hat{C}\left(\mu_0, \hat{\mu}_j^a\right) \end{aligned} \tag{EC.11}$$

We note that the firm's induced target capacity position $\hat{\mu}_j^a$ maximizes the value-to-go function (see equation (EC.4)), and the firm needs to pay a one-time capacity adjustment cost of $\hat{C}\left(\mu_0, \hat{\mu}_j^a\right)$. After the adjustment, the firm does not have another opportunity to change the capacity (recall that $\mathcal{A}(1) = \{0\}$). Therefore, the firm's expected operating profit in period k is simply $H_k(\boldsymbol{\pi}_j, \hat{\mu}_j^a, 1, 0)$, which in turn equals $h_k(\boldsymbol{\pi}_j, \hat{\mu}_j^a)$ from equation (EC.6).

If the firm has not adjusted the capacity ($v_{j-1} = 0$), and decides to delay decision one more period ($u_j = 0$), then we use the superscript s for "stay put", and have the value-to-go function as

$$L_j^s(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1}) = L_j^s(\boldsymbol{\pi_j}, \mu_0, 0) \triangleq h_{j+l}(\boldsymbol{\pi_j}, \mu_0) + E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}}, \mu_0, 0) | \boldsymbol{\pi_j}\right] \tag{EC.12}$$

By delaying the adjustment, the firm earns a profit based on the starting capacity level in this period. However, it maintains the option to change the capacity in the future, as reflected by the term $E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}}, \mu_0, 0) | \boldsymbol{\pi_j}\right]$.

On the other hand, if the firm already adjusted the capacity before, i.e., $v_{j-1} = 1$, then for $k = j, ..., J-l$, we have $\mathcal{A}(v_{k-1}) = \{0\}$ and $u_k = 0$, and we have

$$L_j^s(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1}) = L_j^s(\boldsymbol{\pi_j}, \mu_0, 1) \triangleq \sum_{k=j}^{J-l} H_k(\boldsymbol{\pi}_j, \hat{\mu}_{j-1}, 1, 0) = \sum_{k=j+l}^{J} h_k(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) \tag{EC.13}$$

To sum up, we have the following value-to-go functions contingent upon whether the capacity has been adjusted or not.

$$V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, 0) = V_j(\boldsymbol{\pi_j}, \mu_0, 0) = \max\left\{L_j^a(\boldsymbol{\pi_j}, \mu_0, 0), L_j^s(\boldsymbol{\pi_j}, \mu_0, 0)\right\} \tag{EC.14}$$
$$V_j(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, 1) = L_j^s(\boldsymbol{\pi_j}, \mu_0, 1) \tag{EC.15}$$

When the maximum in equation (EC.14) is attained by $L_j^a(\boldsymbol{\pi_j}, \mu_0, 0)$, it is optimal to adjust the capacity. Otherwise, the firm should delay the adjustment and continue to observe the demand. For ease of exposition, we suppress the dependence on μ_0 and v_{j-1}, and write $V_j(\boldsymbol{\pi_j}, \mu_0, 0)$, $L_j^a(\boldsymbol{\pi_j}, \mu_0, 0)$ and $L_j^s(\boldsymbol{\pi_j}, \mu_0, 0)$ as $V_j(\boldsymbol{\pi_j})$, $L_j^a(\boldsymbol{\pi_j})$ and $L_j^s(\boldsymbol{\pi_j})$ respectively. Therefore, to characterize the firm's optimal policy to stop observing the demand and adjust the capacity, we only need to compare $L_j^a(\boldsymbol{\pi_j})$ and $L_j^s(\boldsymbol{\pi_j})$. Note that, in the single adjustment case, the problem of choosing "when to adjust" and "by how much" is recast as an optimal stopping time problem.

EC.2. Proofs and additional technical details.

EC.2.1. Capacity investment with a single adjustment opportunity

EC.2.1.1. Proofs

In the proofs we focus on the case where the demand distribution is discrete. When the demand distribution is continuous, similar proofs hold.

Proof of Lemma 1. We first show that for any j, we have $E[\boldsymbol{\Pi_{j+1}}|\boldsymbol{\Pi_j}] = \boldsymbol{\Pi_j}$. From equation (1), we have

$$\begin{aligned} E[\Pi_{j+1,i}|\boldsymbol{\Pi_j}] &= \sum_{d_j=0}^{\infty} \frac{\Pi_{j,i} f_j(d_j|\theta_i)}{\sum_{k=1}^{I} [\Pi_{j,k} f_j(d_j|\theta_k)]} Pr(D_j = d_j|\boldsymbol{\Pi_j}) \\ &= \sum_{d_j=0}^{\infty} \frac{\Pi_{j,i} f_j(d_j|\theta_i)}{\sum_{k=1}^{I} [\Pi_{j,k} f_j(d_j|\theta_k)]} \sum_{k=1}^{I} [\Pi_{j,k} f_j(d_j|\theta_k)] \\ &= \Pi_{j,i} \sum_{d_j=0}^{\infty} f_j(d_j|\theta_i) = \Pi_{j,i} \end{aligned} \tag{EC.16}$$

That is, $E[\boldsymbol{\Pi_{j+1}}|\boldsymbol{\Pi_j}] = \boldsymbol{\Pi_j}$. Then for any $j_1 < j_2$, we have

$$\begin{aligned} E[\boldsymbol{\Pi_{j_2}}|\boldsymbol{\Pi_{j_1}}] &= E[E[\boldsymbol{\Pi_{j_2}}|\boldsymbol{\Pi_{j_2-1}}, \boldsymbol{\Pi_{j_1}}]|\boldsymbol{\Pi_{j_1}}] = E[E[\boldsymbol{\Pi_{j_2}}|\boldsymbol{\Pi_{j_2-1}}]|\boldsymbol{\Pi_{j_1}}] \\ &= E[\boldsymbol{\Pi_{j_2-1}}|\boldsymbol{\Pi_{j_1}}] \end{aligned} \tag{EC.17}$$

Applying the above equations iteratively, we have $E[\boldsymbol{\Pi_{j_2}}|\boldsymbol{\Pi_{j_1}}] = \boldsymbol{\Pi_{j_1}}$. □

Proof of Proposition 1. For ease of exposition we define the expected operating profit from period $j+l$ till the end of horizon minus any capacity cost when the firm adjusts capacity from μ_0 to μ in period j, given the information vector of $\boldsymbol{\pi_j}$ as follows.

$$G_j(\boldsymbol{\pi_j}, \mu) \triangleq \sum_{k=j+l}^{J} h_k(\boldsymbol{\pi_j}, \mu) - \hat{C}(\mu_0, \mu) \tag{EC.18}$$

It is observed that for given μ, we have $G_j(\boldsymbol{\pi_j}, \mu)$ is linear in $\boldsymbol{\pi_j}$. For given $\boldsymbol{\pi_j}$, if $\mathcal{K} = \mathbb{R}_+$ and $\mu \in \mathcal{K}$, we have $G_j(\boldsymbol{\pi_j}, \mu)$ is concave in μ; if $\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$, we define $\Delta G_{j,k}(\boldsymbol{\pi_j})$ as follows:

$$\Delta G_{j,k}(\boldsymbol{\pi_j}) \triangleq \frac{G_j(\boldsymbol{\pi_j}, (k+1)\delta) - G_j(\boldsymbol{\pi_j}, k\delta)}{\delta} \text{ for } k \in \mathbb{Z}_+. \tag{EC.19}$$

For given j and $\boldsymbol{\pi_j}$, we have $\Delta G_{j,k}(\boldsymbol{\pi_j})$ is a decreasing sequence in k. In addition, we have $G_j(\boldsymbol{\pi_j},0)<\infty$ and $\lim_{\mu\to\infty} G_j(\boldsymbol{\pi_j},\mu)=-\infty$. In the following, we use $\boldsymbol{x}^T$ to denote the transpose of $\boldsymbol{x}$.

(i). We prove the convexity by induction. By equation (EC.11) and (EC.12), we have $L^a_{J-l+1}(\boldsymbol{\pi_{J-l+1}})$ and $L^s_{J-l+1}(\boldsymbol{\pi_{J-l+1}})$ are linear and therefore convex in $\boldsymbol{\pi_{J-l+1}}$. As the maximum of convex functions is convex, we have $V_{J-l+1}(\boldsymbol{\pi_{J-l+1}})$ is convex.

For $j<J-l$, assume $L^a_{j+1}(\boldsymbol{\pi_{j+1}})$, $L^s_{j+1}(\boldsymbol{\pi_{j+1}})$ and $V_{j+1}(\boldsymbol{\pi_{j+1}})$ are convex. By equation (EC.4) and (EC.11), we have

$$L^a_j(\boldsymbol{\pi_j})=\sup_k\{G_j(\boldsymbol{\pi_j},k\delta)\} \tag{EC.20}$$

For each k, we have $G_j(\boldsymbol{\pi_j},k\delta)$ is linear in $\boldsymbol{\pi_j}$. As the supremum of convex functions is convex and a positive linear combination of convex functions is convex, we have $L^a_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$.

From the induction hypothesis, we have $V_{j+1}(\boldsymbol{\pi_{j+1}})$ is convex, then we can write $V_{j+1}(\boldsymbol{\pi_{j+1}})=\sup_{k\in K_{j+1}}\{\boldsymbol{a}_k\boldsymbol{\pi}^T_{\boldsymbol{j+1}}+b_k\}$, where K_{j+1} represents an index set, $\boldsymbol{a}_k$ is a constant vector of dimensions $1\times I$, $\boldsymbol{\pi}^T_{\boldsymbol{j+1}}$ stands for the transpose of $\boldsymbol{\pi_{j+1}}$, and b_k is a constant. Then define a $1\times I$ vector $\boldsymbol{e}\triangleq(1,...,1)$ and a $I\times I$ diagonal matrix $P_j(d_j)\triangleq diag(f_j(d_j|\theta_1),...,f_j(d_j|\theta_I))$, and following equation (EC.12), we have

$$\begin{aligned}
L^s_j(\boldsymbol{\pi}_j)&=h_{j+l}(\boldsymbol{\pi}_j,\mu_0)+E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}})|\boldsymbol{\pi_j}\right]=h_{j+l}(\boldsymbol{\pi}_j,\mu_0)+E\left[\sup_{k\in K_{j+1}}\left\{\boldsymbol{a}_k\boldsymbol{\Pi}^T_{j+1}+b_k\right\}\middle|\boldsymbol{\pi_j}\right]\\
&=h_{j+l}(\boldsymbol{\pi}_j,\mu_0)+\sum_{d_j=0}^{\infty}\left[\sup_{k\in K_{j+1}}\left\{\boldsymbol{a}_k\frac{P_j(d_j)\boldsymbol{\pi}^T_{\boldsymbol{j}}}{\boldsymbol{e}P_j(d_j)\boldsymbol{\pi}^T_{\boldsymbol{j}}}+b_k\right\}\right]\boldsymbol{e}P_j(d_j)\boldsymbol{\pi}^T_{\boldsymbol{j}}\\
&=h_{j+l}(\boldsymbol{\pi}_j,\mu_0)+\sum_{d_j=0}^{\infty}\left[\sup_{k\in K_{j+1}}\left\{\tilde{\boldsymbol{a}}_k(d_j)\boldsymbol{\pi}^T_{\boldsymbol{j}}\right\}\right]
\end{aligned} \tag{EC.21}$$

where $\tilde{\boldsymbol{a}}_k(d_j)\triangleq\boldsymbol{a}_kP_j(d_j)+b_k\boldsymbol{e}P_j(d_j)$.

Once again, as the supremum of convex functions is convex and a positive linear combination of convex functions is convex, we have $L^s_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$. It follows that $V_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$.

(ii). We show that $\mathbb{P}_j$ is a convex partition of $\mathcal{P}_j$ by verifying the four conditions in Definition 1.

- Condition (i): By the construction of $\mathbb{P}_j$, we have $\emptyset\notin\mathbb{P}_j$.
- Condition (ii): Let $\bigcup_k\mathcal{P}_{j,k}$ denote the union of all sets in $\mathbb{P}_j$. For any $\boldsymbol{\pi_j}\in\bigcup_k\mathcal{P}_{j,k}$, it is trivial that $\boldsymbol{\pi_j}\in\mathcal{P}_j$. Therefore, we have $\bigcup_k\mathcal{P}_{j,k}\subseteq\mathcal{P}_j$. For any $\boldsymbol{\pi_j}\in\mathcal{P}_j$, we have $\Delta G_{j,k}(\boldsymbol{\pi_j})$ decreases in k. As we have $|G_j(\boldsymbol{\pi_j},0)|<\infty$ and $\lim_{\mu\to\infty}G_j(\boldsymbol{\pi_j},\mu)=-\infty$, there exists a k such that $k\delta=\arg\max_{\mu\in\mathcal{K}}G_j(\boldsymbol{\pi_j},\mu)$. Therefore, we have $\boldsymbol{\pi_j}\in\bigcup_k\mathcal{P}_{j,k}$. It follows that $\mathcal{P}_j\subseteq\bigcup_k\mathcal{P}_{j,k}$. Then we have proved that $\bigcup_k\mathcal{P}_{j,k}=\mathcal{P}_j$.

- Condition (iii): Assume there exist $k_1 < k_2$ such that $\mathcal{P}_{j,k_1} \in \mathbb{P}_j$, $\mathcal{P}_{j,k_2} \in \mathbb{P}_j$, and $\mathcal{P}_{j,k_1} \bigcap \mathcal{P}_{j,k_2} \neq \emptyset$. Then for $\boldsymbol{\pi_j} \in \mathcal{P}_{j,k_1} \bigcap \mathcal{P}_{j,k_2}$, we have $\hat{\mu}^a_j(\boldsymbol{\pi_j}) = k_1\delta$ and $\hat{\mu}^a_j(\boldsymbol{\pi_j}) = k_2\delta$. However, this contradicts the fact that $\hat{\mu}^a_j(\boldsymbol{\pi_j})$ is uniquely-defined.
- Condition (iv): Let $\boldsymbol{\pi_j} \in \mathcal{P}_{j,\hat{k}}$ and $\hat{\boldsymbol{\pi}}_{\boldsymbol{j}} \in \mathcal{P}_{j,\hat{k}}$. We have $\hat{\mu}^a_j(\boldsymbol{\pi_j}) = \hat{k}\delta$ and $\hat{\mu}^a_j(\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = \hat{k}\delta$, and

$$L^a_j(\boldsymbol{\pi_j}) = G_j(\boldsymbol{\pi_j}, \hat{k}\delta), \text{ and } L^a_j(\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = G_j(\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}, \hat{k}\delta) \tag{EC.22}$$

We observe that $G_j(\boldsymbol{\pi_j}, \hat{k}\delta)$ is linear in $\boldsymbol{\pi_j}$. From part (i) we have $L^a_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$. Therefore, for any $\alpha \in (0,1)$, we have

$$L^a_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) \le \alpha L^a_j(\boldsymbol{\pi_j}) + (1-\alpha) L^a_j(\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = G_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}, \hat{k}\delta) \tag{EC.23}$$

By the definition of $L^a_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}})$, we have

$$L^a_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = \sup_k \{G_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}, k\delta)\} \ge G_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}, \hat{k}\delta) \tag{EC.24}$$

By equation (EC.23) and (EC.24), we have

$$L^a_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = G_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}, \hat{k}\delta) \tag{EC.25}$$

which implies that $\hat{\mu}^a_j(\alpha\boldsymbol{\pi_j} + (1-\alpha)\hat{\boldsymbol{\pi}}_{\boldsymbol{j}}) = \hat{k}\delta$.

(iii). Consider $\mathcal{P}_{j,k} \in \mathbb{P}_j$. For $\boldsymbol{\pi_j} \in \mathcal{P}_{j,k}$, we have $L^a_j(\boldsymbol{\pi_j})$ is linear in $\boldsymbol{\pi_j}$, and $L^s_j(\boldsymbol{\pi_j})$ is convex in $\boldsymbol{\pi_j}$. Therefore, the difference $\Delta L_j(\boldsymbol{\pi_j}) \triangleq L^a_j(\boldsymbol{\pi_j}) - L^s_j(\boldsymbol{\pi_j})$ is concave in $\boldsymbol{\pi_j}$. Therefore, if $\Delta L_j(\boldsymbol{\pi_j}) \le 0$ for all $\boldsymbol{\pi}_j \in \mathcal{P}_{j,k}$, we have $S_{j,k} = \emptyset$. Otherwise, define $S_{j,k} \triangleq \{\boldsymbol{\pi_j} : \boldsymbol{\pi_j} \in \mathcal{P}_{j,k}, \Delta L_j(\boldsymbol{\pi_j}) > 0\}$. It follows that $S_{j,k}$ is a convex set and for all $\boldsymbol{\pi_j} \in S_{j,k}$, it is optimal for the firm to stop observing the demand and adjust the capacity.

(iv). By equation (EC.20), we have $L^a_j(\boldsymbol{\pi_j}) = \sup_k \{G_j(\boldsymbol{\pi_j}, k\delta)\}$. As $\Delta G_{j,k}(\boldsymbol{\pi_j})$ (defined in equation (EC.19)) is a decreasing sequence in k for given $\boldsymbol{\pi_j}$ and j, to prove the result, it is sufficient to show that for given k, if $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi'_j}$, then $\Delta G_{j,k}(\boldsymbol{\pi_j}) \le \Delta G_{j,k}(\boldsymbol{\pi'_j})$. We first prove the result for the case where for $i_1 < i_2$ and $\epsilon > 0$, we have $\pi'_{j,i_1} = \pi_{j,i_1} - \epsilon$, $\pi'_{j,i_2} = \pi_{j,i_2} + \epsilon$, and $\pi'_{j,\hat{i}} = \pi_{j,\hat{i}}$ for all $\hat{i} \neq i_1, i_2$.

First, by equation (EC.18) and (3), we have

$$\begin{aligned}\Delta G_{j,k}(\boldsymbol{\pi'_j}) - \Delta G_{j,k}(\boldsymbol{\pi_j}) &= \frac{\sum_{i=j+l}^{J} \left[h_i(\boldsymbol{\pi'_j}, (k+1)\delta) - h_i(\boldsymbol{\pi'_j}, k\delta)\right]}{\delta} - \frac{\sum_{i=j+l}^{J} \left[h_i(\boldsymbol{\pi_j}, (k+1)\delta) - h_i(\boldsymbol{\pi_j}, k\delta)\right]}{\delta} \\ &= \frac{\epsilon}{\delta} \sum_{i=j+l}^{J} \Big\{ E\left[g(D_i)|\theta_{i_2}\right] - E\left[g(D_i)|\theta_{i_1}\right] \Big\} \end{aligned} \tag{EC.26}$$

where

$$g(D_i) \triangleq -c_1[D_i - (k+1)\delta\tau]^+ + c_1(D_i - k\delta\tau)^+ - c_0\delta\tau.$$

Because $D_i|\theta_{i_1} \preceq_{st} D_i|\theta_{i_2}$, and $g(D_i)$ increases in D_i, we have

$$\Delta G_{j,k}(\boldsymbol{\pi'_j}) - \Delta G_{j,k}(\boldsymbol{\pi_j}) \geq 0. \tag{EC.27}$$

For an arbitrary pair of $\boldsymbol{\pi_j}$ and $\boldsymbol{\pi'_j}$ such that $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi'_j}$, we observe that $\boldsymbol{\pi'_j}$ can be obtained from $\boldsymbol{\pi_j}$ within finite steps using the operations above (subtract $\epsilon_{j,i}$ from an element with a lower index and add $\epsilon_{j,i}$ to an element with a higher index). □

Proof of Proposition 2. The proof of part (i) is similar to the proof for Proposition 1(i). We only prove part (ii) here.

When the feasible set for the capacity adjustment is continuous, we have

$$L_j^a(\boldsymbol{\pi_j}) = \max_{\mu \in \mathbb{R}_+} \left\{ G_j(\boldsymbol{\pi_j}, \mu) \right\}. \tag{EC.28}$$

We first define

$$\hat{G}_j(\boldsymbol{\pi_j}, \boldsymbol{\pi'_j}, \mu) \triangleq G_j(\boldsymbol{\pi'_j}, \mu) - G_j(\boldsymbol{\pi_j}, \mu). \tag{EC.29}$$

Then to prove the result, it is sufficient to show that for given μ, if $\boldsymbol{\pi_j} \preceq \boldsymbol{\pi'_j}$, we have $\frac{\partial \hat{G}_j}{\partial \mu}(\boldsymbol{\pi_j}, \boldsymbol{\pi'_j}, \mu) \geq 0$. Following a similar step as in the proof of Proposition 1(iv), it is sufficient to prove for the following case: for $i_1 < i_2$ and $\epsilon > 0$, we have $\pi_{j',i_1} = \pi_{j,i_1} - \epsilon$, $\pi_{j',i_2} = \pi_{j,i_2} + \epsilon$, and $\pi_{j',\hat{i}} = \pi_{j,\hat{i}}$ for all $\hat{i} \neq i_1, i_2$.

Following equation (EC.18) and (3), we have

$$\begin{aligned} \frac{\partial \hat{G}_j}{\partial \mu}(\boldsymbol{\pi_j}, \boldsymbol{\pi'_j}, \mu) &= \sum_{i=j+l}^{J} \left[\frac{\partial h_i}{\partial \mu}(\boldsymbol{\pi'_j}, \mu) - \frac{\partial h_i}{\partial \mu}(\boldsymbol{\pi_j}, \mu) \right] \\ &= \epsilon c_1 \tau \sum_{i=j+l}^{J} \left[F_i(\mu\tau|\theta_{i_1}) - F_i(\mu\tau|\theta_{i_2}) \right] \geq 0 \end{aligned} \tag{EC.30}$$

It follows that $\hat{\mu}_j^a(\boldsymbol{\pi_j}) \leq \hat{\mu}_j^a(\boldsymbol{\pi'_j})$, which completes the proof. □

We state the following proposition from Gallego (1992) before proving Proposition 3.

Proposition EC.1 **(Proposition 1 in Gallego (1992))**. *Let $\mathcal{F}$ denote the class of cumulative distributions with finite mean μ and variance σ^2, and R be a finite constant.*

$$\max_{F \in \mathcal{F}} \int (x - R)^+ dF(x) = \frac{1}{2}(\sqrt{\Delta^2 + \sigma^2} - \Delta) \tag{EC.31}$$

where $\Delta = R - \mu$.

Essentially, this is a one-sided deviation bound. Following a similar proof, we have

$$\max_{F \in \mathcal{F}} \int (R - x)^+ dF(x) = \frac{1}{2}(\sqrt{\Delta^2 + \sigma^2} - \Delta) \tag{EC.32}$$

where $\Delta = \mu - R$.

Proof of Proposition 3. We derive an upper bound of the regret as follows. We first observe that for any x, y, and z, we have

$$(x-y)^+ \le (x-z)^+ + (z-y)^+. \tag{EC.33}$$

We have $\hat{\lambda}_{i,\tau_n} = \frac{N(n\lambda_i\tau_n)}{n\tau_n}$ from equation (13). To simplify the notations, we similarly define

$$\hat{\lambda}_{i,j\tau_n} \triangleq \frac{D_j|\theta_{i,n}}{n\tau_n} = \frac{N(n\lambda_i j\tau_n) - N(n\lambda_i(j-1)\tau_n)}{n\tau_n} \text{ for } j = 2,3,...,J_n \tag{EC.34}$$

We observe that $\hat{\lambda}_{i,j\tau_n}$ for $j=1,2,...,J_n$ is a sequence of i.i.d. random variables with $E\left(\hat{\lambda}_{i,j\tau_n}\right) = \lambda_i$ and $\text{Var}\left(\hat{\lambda}_{i,j\tau_n}\right) = \frac{\sigma^2\lambda_i}{n\tau_n}$. From equation (14), we have

$$V_{0,n}^{ts}(\boldsymbol{\pi_1}) = \sum_{i=1}^{I} \pi_{1,i} E\left\{ \begin{array}{l} pn\hat{\lambda}_{i,(l_n+1)\tau_n}\tau_n - c_1 n\left(\hat{\lambda}_{i,(l_n+1)\tau_n} - \mu_0\right)^+ \tau_n - c_0 n\mu_0\tau_n - \hat{C}(n\mu_0, n\hat{\lambda}_{i,\tau_n}) \\ + \sum_{j=l_n+2}^{J_n} \left[pn\hat{\lambda}_{i,j\tau_n}\tau_n - c_1 n\left(\hat{\lambda}_{i,j\tau_n} - \hat{\lambda}_{i,\tau_n}\right)^+ \tau_n - c_0 n\hat{\lambda}_{i,\tau_n}\tau_n \right] \end{array} \middle| \theta_{i,n} \right\} \tag{EC.35}$$

We show analysis for the case where $\lambda_i \ge \mu_0$ for all i, as the analysis for the other case is similar. The expected operating profit from the period l_n+1, $pn\lambda_i\tau_n - c_1 nE\left(\hat{\lambda}_{i,(l_n+1)\tau_n} - \mu_0\right)^+ \tau_n - c_0 n\mu_0\tau_n$, is positive, as $p \ge c_1 > c_0$. We also note that the lead-time in the unit of time l_t satisfies that $l_t = l_n\tau_n$.Therefore, we have

$$\text{RHS of (EC.35)} \ge \sum_{i=1}^{I} \pi_{1,i} \left\{ \begin{array}{l} (p-c_0)n\lambda_i(T-\tau_n-l_t) - E\left[\hat{C}(n\mu_0, n\hat{\lambda}_{i,\tau_n})\right] \\ -c_1 n\tau_n \sum_{j=l_n+2}^{J_n} E\left(\hat{\lambda}_{i,j\tau_n} - \hat{\lambda}_{i,\tau_n}\right)^+ \end{array} \right\} \tag{EC.36}$$

By equation (16), we have the deterministic upper bound

$$V_{0,n}^d = \sum_{i=1}^{I} \pi_{1,i} \Big\{ (p-c_0)n\lambda_i(T-l_t) - \hat{C}(n\mu_0, n\lambda_i) \Big\}.$$

Therefore, combining equation (EC.36) and the expression above for $V_{0,n}^d$, we have the regret

$$\begin{aligned} R_n^{ts} &= 1 - V_{0,n}^{ts}/V_{0,n}^d \\ &\le \frac{1}{V_{0,n}^d} \sum_{i=1}^{I} \pi_{1,i} \left\{ (p-c_0)n\lambda_i\tau_n - \hat{C}(n\mu_0, n\lambda_i) + E\left[\hat{C}(n\mu_0, n\hat{\lambda}_{i,\tau_n})\right] + c_1 n\tau_n \sum_{j=l_n+2}^{J_n} E\left(\hat{\lambda}_{i,j\tau_n} - \hat{\lambda}_{i,\tau_n}\right)^+ \right\}. \end{aligned} \tag{EC.37}$$

Recall that for the initial capacity position μ and target capacity position μ', we have $\hat{C}(\mu,\mu') = c_a(\mu'-\mu)^+ + \gamma_a(\mu-\mu')^+$. We notice when $\gamma_a \ge 0$, we can directly apply (EC.33) and $E[-\gamma_a(n\mu_0 - n\lambda_i)^+ + \gamma_a(n\mu_0 - n\hat{\lambda}_{i,\tau_n})^+] \le E[\gamma_a n(\lambda_i - \hat{\lambda}_{i,\tau_n})^+]$. When $\gamma_a < 0$, we have $E[-\gamma_a(n\mu_0 - n\lambda_i)^+ + \gamma_a(n\mu_0 - n\hat{\lambda}_{i,\tau_n})^+] \le 0$ by Jensen's inequality. Following the discussions, we have established that

$-\hat{C}(n\mu_0, n\lambda_i) + E\left[\hat{C}(n\mu_0, n\hat{\lambda}_{i,\tau_n})\right] \le c_a nE(\hat{\lambda}_{i,\tau_n} - \lambda_i)^+ + \gamma_a^+ nE(\lambda_i - \hat{\lambda}_{i,\tau_n})^+$. Therefore, applying (EC.33), we have

$$\text{RHS of (EC.37)} \le \frac{1}{V_{0,n}^d}\sum_{i=1}^{I}\pi_{1,i}\left\{\begin{array}{l}(p-c_0)n\lambda_i\tau_n + c_a nE(\hat{\lambda}_{i,\tau_n}-\lambda_i)^+ + \gamma_a^+ nE(\lambda_i - \hat{\lambda}_{i,\tau_n})^+ \\ +c_1 n\tau_n \sum_{j=l_n+2}^{J_n}\left[E\left(\hat{\lambda}_{i,j\tau_n}-\lambda_i\right)^+ + E\left(\lambda_i - \hat{\lambda}_{i,\tau_n}\right)^+\right]\end{array}\right\} \tag{EC.38}$$

From equation (EC.38), it is clear that to derive an upper bound of R_n^{ts}, we need to find upper bounds for $E\left(\hat{\lambda}_{i,j\tau_n}-\lambda_i\right)^+$ and $E\left(\lambda_i-\hat{\lambda}_{i,\tau_n}\right)^+$ respectively. Recall that $E\left(\hat{\lambda}_{i,j\tau_n}\right)=\lambda_i$ and $\text{Var}\left(\hat{\lambda}_{i,j\tau_n}\right)=\frac{\sigma^2\lambda_i}{n\tau_n}$. In the following, we use C_i to represent a constant for all i, which is independent of n and τ_n.

We first find an upper bound for $E\left(\hat{\lambda}_{i,j\tau_n}-\lambda_i\right)^+$. By equation (EC.31), we have

$$E\left(\hat{\lambda}_{i,j\tau_n}-\lambda_i\right)^+ \le \frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n\tau_n}} \text{ for } j=1,2,...,J_n. \tag{EC.39}$$

For $E\left(\lambda_i-\hat{\lambda}_{i,\tau_n}\right)^+$, by equation (EC.32), we have the following

$$E\left(\lambda_i-\hat{\lambda}_{i,\tau_n}\right)^+ \le \frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n\tau_n}}. \tag{EC.40}$$

From equation (EC.39) and (EC.40), we have

$$\text{RHS of (EC.38)} \le C_1\tau_n + \frac{C_2}{\sqrt{n\tau_n}} \tag{EC.41}$$

Then the result follows by setting that $\tau_n \asymp n^{-\frac{1}{3}}$. □

EC.2.1.2. Remarks

In this section, we briefly discuss some of the modeling features and assumptions, the rationale behind them, and the consequences of removing or relaxing them.

General demand process: The optimal policy characterized in Section 4.1 can be applied to a large class of random variables and demand processes. In our base model, we have finite demand types and each type is characterized by a demand type parameter. However, our model can be extended to accommodate more general features. First, demand type i can be characterized by a vector of parameters $\boldsymbol{\theta_i}$. We only require the demand stochastically increases in the demand type index, i.e., $D_j|\boldsymbol{\theta_{i_1}} \preceq_{st} D_j|\boldsymbol{\theta_{i_2}}$ for $i_1 \le i_2$. Thus as long as the demand type forms an ordered set, our results apply. Second, if there are uncountably infinite demand types, i.e., the prior and posterior distributions are characterized by a continuous distribution function, Proposition 1 and 2 still hold. That is, assuming the firm decides to adjust the capacity, the target capacity increases as the likelihood of demand being high increases. As in the base model, the decision to adjust the capacity is not monotone in the likelihood. Finally, the optimal policy still holds when the demand is non-stationary; for example, $D_j|\theta_i$ may represent a non-stationary Poisson process with the mean demand $\lambda_j(\theta_i)$ following a Bass diffusion curve where the market size is θ_i and the coefficient of innovation and coefficient of imitation are fixed across all the demand types. In this case, the random term $\xi_j|\theta_i$ represents a "shifted Poisson" distribution, which has mean 0 and variance $\lambda_j(\theta_i)$.

Censored demand: As the demand beyond capacity is satisfied by an outside option (e.g., outsourcing, overtime, or temporarily using the capacity designated for a different product), demand is fully observed and not censored. However, our model can be extended to accommodate censored demand. In the case of unobservable lost sales, the posterior distribution can be updated as follows:

$$\pi_{j+1,i} = \begin{cases} \frac{\pi_{j,i} f_j(d_j|\theta_i)}{\sum_{k=1}^{I}\left[\pi_{j,k} f_j(d_j|\theta_k)\right]} & \text{if } d_j < \mu \\ \frac{\pi_{j,i} Pr(D_j \geq \mu|\Theta=\theta_i)}{\sum_{k=1}^{I}\left[\pi_{j,k} Pr(D_j \geq \mu|\Theta=\theta_k)\right]} & \text{if } d_j \geq \mu \end{cases} \tag{EC.42}$$

Correspondingly, the firm's expected profit in one period is changed as follows.

$$h_j^c(\boldsymbol{\pi_j}, \mu) = \sum_{i=1}^{I} \pi_{j,i} E\left[p \min\{D_j, \mu\tau\} - c_0\mu\tau|\Theta = \theta_i\right] \tag{EC.43}$$

Then, following a similar process of defining equation (EC.10), we can define the value-to-go function $V_j^c(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, v_{j-1})$. Following a similar proof as that of Proposition 1, we can show that an optimal policy with similar structure holds.

Fixed cost: If the adjustment decision is associated with a fixed cost, the cost associated with changing the capacity level from μ to μ', denoted by $\check{C}(\mu, \mu')$, is

$$\check{C}(\mu, \mu') \triangleq c_a(\mu' - \mu)^+ + \overline{K}\mathbb{1}_{\{\mu'-\mu>0\}} + \gamma_a(\mu - \mu')^+ + \underline{K}\mathbb{1}_{\{\mu-\mu'>0\}}, \tag{EC.44}$$

where $\overline{K}$ is the fixed cost associated with capacity investment, and $\underline{K}$ is the fixed cost associated with capacity disinvestment. In this case, the optimal policy remains the same. However, as it becomes more costly to adjust capacity, the adjustment region shrinks when the fixed cost increases.

Leadtime: Our results can be generalized to the case where the leadtime to invest in capacity (say, $\bar{l}$) is different from the leadtime to disinvest, $\underline{l}$. To see this, first consider the case where $\bar{l} > \underline{l}$. In period j, if the firm decides to invest, the invested capacity will be available in period $j + \bar{l}$. Therefore, the target capacity position from equation (EC.4) should be modified as

$$\overline{\mu}_j^a(\boldsymbol{\pi_j}) = \arg\max_{\mu \in \mathcal{K}, \mu \geq \mu_0} \left\{ \sum_{k=j+\bar{l}}^{J} h_k(\boldsymbol{\pi_j}, \mu) - c_a\mu \right\}. \tag{EC.45}$$

Similarly, if the firm decides to disinvest, the target capacity position should be modified as

$$\underline{\mu}_j^a(\boldsymbol{\pi_j}) = \underset{\mu\in\mathcal{K},0\le\mu\le\mu_0}{\arg\max}\left\{\sum_{k=j+\underline{l}}^{J} h_k(\boldsymbol{\pi_j},\mu)+\gamma_a\mu\right\}. \tag{EC.46}$$

Then we define the firm's value-to-go function $\overline{L}_j^a(\boldsymbol{\pi_j})$ when the firm invests to $\overline{\mu}_j^a(\boldsymbol{\pi_j})$ and the value-to-go function $\underline{L}_j^a(\boldsymbol{\pi_j})$ when the firm disinvests to $\underline{\mu}_j^a(\boldsymbol{\pi_j})$ as follows.

$$\overline{L}_j^a(\boldsymbol{\pi_j}) = \sum_{k=j+\underline{l}}^{j+\bar{l}-1} h_k(\boldsymbol{\pi_j},\mu_0) + \sum_{k=j+\bar{l}}^{J} h_k\left(\boldsymbol{\pi_j},\overline{\mu}_j^a(\boldsymbol{\pi_j})\right) - c_a\left(\overline{\mu}_j^a(\boldsymbol{\pi_j})-\mu_0\right)$$

$$\underline{L}_j^a(\boldsymbol{\pi_j}) = \sum_{k=j+\underline{l}}^{J} h_k\left(\boldsymbol{\pi_j},\underline{\mu}_j^a(\boldsymbol{\pi_j})\right) - \gamma_a\left(\mu_0-\underline{\mu}_j^a(\boldsymbol{\pi_j})\right) \tag{EC.47}$$

Then, we define the value-to-go function when the firm adjusts capacity and the associated target capacity position as

$$L_j^a(\boldsymbol{\pi_j}) = \max\{\overline{L}_j^a(\boldsymbol{\pi_j}),\underline{L}_j^a(\boldsymbol{\pi_j})\} \text{ and } \hat{\mu}_j^a(\boldsymbol{\pi_j}) = \begin{cases}\overline{\mu}_j^a(\boldsymbol{\pi_j}) & \text{if } L_j^a(\boldsymbol{\pi_j}) = \overline{L}_j^a(\boldsymbol{\pi_j})\\ \underline{\mu}_j^a(\boldsymbol{\pi_j}) & \text{otherwise}\end{cases} \tag{EC.48}$$

When the firm chooses to stay put, the value-to-go function is as follows:

$$L_j^s(\boldsymbol{\pi_j}) = h_{j+\underline{l}}(\boldsymbol{\pi_j},\mu_0) + E\left[V_{j+1}(\boldsymbol{\Pi_{j+1}})|\,\boldsymbol{\pi_j}\right] \tag{EC.49}$$

Then following similar proof of Proposition 1 and 2, the same structure of the optimal policy holds.

Alternative demand process in the two-step heuristic: Our results and analysis also hold for a stationary demand process where the demand in each period, $\hat{D}_j|\theta_{i,n}$, is a sequence of i.i.d. random variables, whose mean is $n\lambda_i\tau_n$ and variance is $n^{\delta_1}\sigma^2\tau_n$. The following corollary shows that our heuristic is asymptotically optimal with a convergence rate that depends on the variance term, assuming the variance does not increase too fast with respect to the problem size n, i.e., $\delta_1 < 2$.

COROLLARY EC.1. *For a stationary demand sequence $\hat{D}_j|\theta_{i,n}$ whose mean is $n\lambda_i\tau_n$ and variance is $n^{\delta_1}\sigma^2\tau_n$, if $\delta_1 < 2$ and $\tau_n \asymp n^{\frac{\delta_1-2}{3}}$ for all n, the two-step heuristic is asymptotically optimal and $R_n^{ts} = O\left(n^{\frac{\delta_1-2}{3}}\right)$.*

Intuitively, as the variance term increases with respect to the size of the problem, the firm is able to extract less information per unit time. Therefore, it takes the firm more time to learn about demand to justify a capacity adjustment.

Proof of Corollary EC.1 The proof is similar to the proof of Proposition 3. Therefore, we only show the key steps. We derive an upper bound of the regret as follows. To simplify the notations, we define

$$\check{\lambda}_{i,j\tau_n} \triangleq \frac{\hat{D}_j|\theta_{i,n}}{n\tau_n} \text{ for } j = 1,2,...,J_n \tag{EC.50}$$

We observe that $\check{\lambda}_{i,j\tau_n}$ for $j=1,2,...,J_n$ is a sequence of i.i.d. random variables with $E\left(\check{\lambda}_{i,j\tau_n}\right)=\lambda_i$ and $\mathrm{Var}\left(\check{\lambda}_{i,j\tau_n}\right)=\frac{\sigma^2}{n^{2-\delta_1}\tau_n}$. From equation (14), we have

$$V_{0,n}^{ts}(\boldsymbol{\pi_1})=\sum_{i=1}^{I}\pi_{1,i}E\left\{\begin{array}{l} pn\check{\lambda}_{i,(l_n+1)\tau_n}\tau_n-c_1n\left(\check{\lambda}_{i,(l_n+1)\tau_n}-\mu_0\right)^+\tau_n-c_0n\mu_0\tau_n-\hat{C}(n\mu_0,n\check{\lambda}_{i,\tau_n}) \\ +\sum_{j=l_n+2}^{J_n}\left[pn\check{\lambda}_{i,j\tau_n}\tau_n-c_1n\left(\check{\lambda}_{i,j\tau_n}-\check{\lambda}_{i,\tau_n}\right)^+\tau_n-c_0n\check{\lambda}_{i,\tau_n}\tau_n\right]\end{array}\middle|\theta_{i,n}\right\} \tag{EC.51}$$

We show analysis for the case where $\lambda_i\geq\mu_0$ for all i, as the analysis for the other case is similar. The expected operating profit from the period l_n+1, $pn\lambda_i\tau_n-c_1nE\left(\check{\lambda}_{i,(l_n+1)\tau_n}-\mu_0\right)^+\tau_n-c_0n\mu_0\tau_n$, is positive, as $p\geq c_1>c_0$. Therefore, we have

$$\text{RHS of (EC.51)}\geq\sum_{i=1}^{I}\pi_{1,i}\left\{\begin{array}{l}(p-c_0)n\lambda_i(T-\tau_n-l_t)-E\left[\hat{C}(n\mu_0,n\check{\lambda}_{i,\tau_n})\right] \\ -c_1n\tau_n\sum_{j=l_n+2}^{J_n}E\left(\check{\lambda}_{i,j\tau_n}-\check{\lambda}_{i,\tau_n}\right)^+\end{array}\right\} \tag{EC.52}$$

By equation (16), we still have the deterministic upper bound

$$V_{0,n}^d=\sum_{i=1}^{I}\pi_{1,i}\Big\{(p-c_0)n\lambda_i(T-l_t)-\hat{C}(n\mu_0,n\lambda_i)\Big\}.$$

Therefore, combining equation (EC.52) and the expression above for $V_{0,n}^d$, we have the regret

$$\begin{aligned}R_n^{ts}&=1-V_{0,n}^{ts}/V_{0,n}^d\\ &\leq\frac{1}{V_{0,n}^d}\sum_{i=1}^{I}\pi_{1,i}\left\{(p-c_0)n\lambda_i\tau_n-\hat{C}(n\mu_0,n\lambda_i)+E\left[\hat{C}(n\mu_0,n\hat{\lambda}_{i,\tau_n})\right]+c_1n\tau_n\sum_{j=l_n+2}^{J_n}E\left(\check{\lambda}_{i,j\tau_n}-\check{\lambda}_{i,\tau_n}\right)^+\right\}.\end{aligned} \tag{EC.53}$$

Recall that for the initial capacity position μ and target capacity position μ', we have $\hat{C}(\mu,\mu')=c_a(\mu'-\mu)^++\gamma_a(\mu-\mu')^+$. We notice when $\gamma_a\geq0$, we can directly apply (EC.33) and $E[-\gamma_a(n\mu_0-n\lambda_i)^++\gamma_a(n\mu_0-n\check{\lambda}_{i,\tau_n})^+]\leq E[\gamma_an(\lambda_i-\check{\lambda}_{i,\tau_n})^+]$. When $\gamma_a<0$, we have $E[-\gamma_a(n\mu_0-n\lambda_i)^++\gamma_a(n\mu_0-n\check{\lambda}_{i,\tau_n})^+]\leq0$ by Jensen's inequality. Therefore, applying (EC.33), we have

$$\text{RHS of (EC.53)}\leq\frac{1}{V_{0,n}^d}\sum_{i=1}^{I}\pi_{1,i}\left\{\begin{array}{l}(p-c_0)n\lambda_i\tau_n+c_anE(\check{\lambda}_{i,\tau_n}-\lambda_i)^++\gamma_a^+nE(\lambda_i-\check{\lambda}_{i,\tau_n})^+ \\ +c_1n\tau_n\sum_{j=l_n+2}^{J_n}\left[E\left(\check{\lambda}_{i,j\tau_n}-\lambda_i\right)^++E\left(\lambda_i-\check{\lambda}_{i,\tau_n}\right)^+\right]\end{array}\right\} \tag{EC.54}$$

From equation (EC.54), it is clear that to derive an upper bound of R_n^{ts}, we need to find upper bounds for $E\left(\check{\lambda}_{i,j\tau_n}-\lambda_i\right)^+$ and $E\left(\lambda_i-\check{\lambda}_{i,\tau_n}\right)^+$ respectively. Recall that $E\left(\check{\lambda}_{i,j\tau_n}\right)=\lambda_i$ and $\mathrm{Var}\left(\check{\lambda}_{i,j\tau_n}\right)=\frac{\sigma^2}{n^{2-\delta_1}\tau_n}$. In the following, we use $\check{C}_i$ to represent a constant for all i, which is independent of n and τ_n.

We first find an upper bound for $E\left(\check{\lambda}_{i,j\tau_n}-\lambda_i\right)^+$. By equation (EC.31), we have

$$E\left(\check{\lambda}_{i,j\tau_n}-\lambda_i\right)^+\leq\frac{\sigma}{2n\sqrt{\tau_n}}\text{ for }j=1,2,...,J_n. \tag{EC.55}$$

For $E\left(\lambda_i - \check{\lambda}_{i,\tau_n}\right)^+$, by equation (EC.32), we have the following

$$E\left(\lambda_i - \check{\lambda}_{i,\tau_n}\right)^+ \leq \frac{\sigma}{2\sqrt{n^{2-\delta_1}\tau_n}}. \tag{EC.56}$$

From equation (EC.55) and (EC.56), we have

$$\text{RHS of (EC.54)} \leq \check{C}_1 \tau_n + \frac{\check{C}_2}{\sqrt{n^{2-\delta_1}\tau_n}} \tag{EC.57}$$

Then the result follows by setting that $\tau_n \asymp n^{\frac{\delta_1 - 2}{3}}$. □

EC.2.2. Capacity investment with multiple adjustment opportunities

EC.2.2.1. Proofs

Without loss of generality, we consider $\mathcal{K} = \mathbb{Z}_+$ in the following proof. In the general case of $\mathcal{K} = \{k\delta : k \in \mathbb{Z}_+\}$, the following proof holds by defining new decision and state variables $\check{\mu} \triangleq \hat{\mu}/\delta$ and $\check{\mu}_j \triangleq \hat{\mu}_j/\delta$.

Lemma EC.1 (**$L^\natural$-concavity preservation**). *Let $S \subset \mathbb{Z}_+$ be a lattice.*

(i) If $f(\boldsymbol{v})$ and $g(\boldsymbol{v})$ are $L^\natural$-concave in $\boldsymbol{v}$, then $\alpha f(\boldsymbol{v}) + \beta g(\boldsymbol{v})$ where $\alpha, \beta \geq 0$ is also $L^\natural$-concave.

(ii) If $f(\boldsymbol{v}, \xi)$ is $L^\natural$-concave, then $g(\boldsymbol{v}) = \max_{\xi \in S} f(\boldsymbol{v}, \xi)$ is also $L^\natural$-concave.

Proof of Lemma EC.1 We note the part (i) is trivial. We prove part (ii) below following a similar approach of Lemma 2 in Zipkin (2008).

If $f(\boldsymbol{v}, \xi)$ is $L^\natural$-concave, then we have $r(\boldsymbol{v}, \xi, \zeta) = f[(\boldsymbol{v}, \xi) - \zeta(\boldsymbol{e}, 1)]$ is supermodular. Therefore, we have

$$\begin{aligned}\psi(\boldsymbol{v}, \zeta) = g(\boldsymbol{v} - \zeta\boldsymbol{e}) = \max_{\xi \in S} f(\boldsymbol{v} - \zeta\boldsymbol{e}, \xi) = \max_{\xi \in S} f[(\boldsymbol{v}, \xi + \zeta) - \zeta(\boldsymbol{e}, 1)] = \max_{\xi \in S} r(\boldsymbol{v}, \xi + \zeta, \zeta) \\ = \max_{\epsilon \geq \zeta \& \epsilon - \zeta \in S} r(\boldsymbol{v}, \epsilon, \zeta)\end{aligned}$$

Note the set $\{(\epsilon, \zeta) : \epsilon \geq \zeta, \epsilon, \zeta \in S\}$ is a sublattice of $\mathbb{Z}_+^2$. Therefore, the maximum over ϵ is supermodular following Theorem 2.7.6 in Topkis (1998). As a result, we have $g(\boldsymbol{v})$ is $L^\natural$-concave. □

To proceed, we use the following definition of Lovász extension on the hypercube $\{L, U\}^n$ (Lovász 1983, Chen et al. 2014) to obtain a continuous extension of an $L^\natural$-concave function. Here, the hypercube $\{L, U\}^n$ refers to the set $\{\boldsymbol{z} \in \mathbb{Z}_+^n : z_j \in \{L, U\}, j = 1, 2, ..., n\}$.

Definition EC.1 (Lovász extension). Given a discrete function $f : \{L, U\}^n \to \mathbb{R}$, for any given point $\boldsymbol{x} \in [L, U]^n$, let σ be the permutation of $\{1, 2, ..., n\}$ such that $x_{\sigma(1)} \geq x_{\sigma(2)} \geq ... \geq x_{\sigma(n)}$. For $0 \leq i \leq n$, define $\boldsymbol{y}(i) \in \{L, U\}^n$ such that $y_{\sigma(1)}(i) = y_{\sigma(2)}(i) = ... = y_{\sigma(i)}(i) = U$ and $y_{\sigma(i+1)}(i) = y_{\sigma(i+2)}(i) = ... = y_{\sigma(n)}(i) = L$. Let λ_i be the unique coefficient of $\boldsymbol{y}(i), i = 0, 1, ..., n$; i.e.,

$$\boldsymbol{x} = \sum_{i=0}^{n} \lambda_i \boldsymbol{y}(i) \tag{EC.58}$$

The Lovász extension $f^L:[L,U]^n\to\mathbb{R}$ of function f at point $\boldsymbol{x}$ is defined as follows:

$$f^L(\boldsymbol{x})=\sum_{i=0}^{n}\lambda_i f(\boldsymbol{y}(i)). \tag{EC.59}$$

The following results are due to Lovász (1983) and Murota (2003). We skip the proof of the lemma and refer interested readers to Murota (2003) and Chen et al. (2014).

Lemma EC.2 **(Concave extension of an $L^\natural$-concave function)**.

(i) For a discrete function f defined on a hypercube $\{L,U\}^n$, its concave envelope is identical to its Lovász extension if and only if f is supermodular.

(ii) For a discrete $L^\natural$-concave function f, its global extension can be obtained by first obtaining its Lovász extension for every unit hypercube in its domain, and then paste all these extensions together.

Lemma EC.3 **($L^\natural$-concavity of value-to-go functions)**. *For all $j\in\{1,2,...,J-l+1\}$,*

(i) $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)$ is $L^\natural$-concave in $(\hat\mu_{j-1},\hat\mu)$.

(ii) $V_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1})$ is $L^\natural$-concave in $\hat\mu_{j-1}$.

Proof of Lemma EC.3 We first prove part (i). By definition, in order to show $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)$ is $L^\natural$-concave in $(\hat\mu_{j-1},\hat\mu)$, we need to show that $\psi_j^h(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu,\zeta)=H_j^m[(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)-\zeta(\mathbf{0},1,1)]=H_j^m[(\boldsymbol{\pi_j},\hat\mu_{j-1}-\zeta,\hat\mu-\zeta)$ is supermodular in $(\hat\mu_{j-1},\hat\mu,\zeta)$. For $j=J-l+1$, it is trivial as $H_{J-l+1}^m(\boldsymbol{\pi_{J-l+1}},\hat\mu_{J-l}-\zeta,\hat\mu-\zeta)=0$. For any $j=1,2,...,J-l$, we have $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1}-\zeta,\hat\mu-\zeta)=h_{j+l}(\boldsymbol{\pi_j},\hat\mu-\zeta)-\hat C(\hat\mu_{j-1},\hat\mu)$. Note that $h_{j+l}(\boldsymbol{\pi_j},\hat\mu-\zeta)$ is supermodular in $(\hat\mu,\zeta)$ and $-\hat C(\hat\mu_{j-1},\hat\mu)$ is supermodular in $(\hat\mu,\hat\mu_{j-1})$, and we have $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1}-\zeta,\hat\mu-\zeta)$ is supermodular in $(\hat\mu_{j-1},\hat\mu,\zeta)$. Therefore, we have proved that $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)$ is $L^\natural$-concave in $(\hat\mu_{j-1},\hat\mu)$.

We next prove part (ii) of this lemma by induction. For period $J-l+1$, we have $V_{J-l+1}^m(\boldsymbol{\pi_{J-l+1}},\hat\mu_{J-l})=0$. Therefore, it is trivial that $V_{J-l+1}^m(\boldsymbol{\pi_{J-l+1}},\hat\mu_{J-l})$ is $L^\natural$-concave in $\hat\mu_{J-l}$.

Assume that for period $j+1$, we have that $V_{j+1}^m(\boldsymbol{\pi_{j+1}},\hat\mu_j)$ is $L^\natural$-concave in $\hat\mu_j$. For period j, we have $V_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1})=\max_{\hat\mu\in\mathcal{K}}\{H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)+E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}},\hat\mu)|\boldsymbol{\pi_j}\right]\}$. By the induction hypothesis and Lemma EC.1 (i), we have $E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}},\hat\mu)|\boldsymbol{\pi_j}\right]$ is $L^\natural$-concave in $\hat\mu_j$ and therefore it is also $L^\natural$-concave in $(\hat\mu_{j-1},\hat\mu)$. Note that we have proved that $H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)$ is $L^\natural$-concave in $(\hat\mu_{j-1},\hat\mu)$. Then by Lemma EC.1 (i) and (ii) and the fact that $\mathcal{K}$ is also a lattice, we have $V_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1})$ is $L^\natural$-concave in $\hat\mu_{j-1}$. $\square$

Proof of Proposition 4 For any $j=1,2,...,J-l+1$ and any $\hat\mu\in\mathcal{K}$, we define

$$f(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)\triangleq H_j^m(\boldsymbol{\pi_j},\hat\mu_{j-1},\hat\mu)+E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}},\hat\mu)|\boldsymbol{\pi_j}\right] \tag{EC.60}$$

Correspondingly, we denote the Lovász extension of f by f^L. By Lemma EC.2 and the definition of Lovász extension, we note that f^L is a piecewise linear concave function and we have the following:

$$V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = \max_{\hat{\mu}\in\mathcal{K}} f(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) = \max_{\hat{\mu}\in\mathbb{R}_+} f^L(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) \tag{EC.61}$$

We next show the optimal policy follows the control band structure. For $\hat{\mu}_{j-1} = 0$, we have $f(\boldsymbol{\pi_j}, 0, \hat{\mu}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a\hat{\mu} + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]$ and the corresponding Lovász extension by $f^L(\boldsymbol{\pi_j}, 0, \hat{\mu})$. Then we define

$$\underline{\mu}_j(\boldsymbol{\pi_j}) = \arg\max_{\hat{\mu}\in\mathbb{R}_+} \left\{f^L(\boldsymbol{\pi_j}, 0, \hat{\mu})\right\} \tag{EC.62}$$

As f^L is the piecewise linear concave envelope for f, we have $\underline{\mu}_j(\boldsymbol{\pi_j}) \in \mathcal{K}$. For any $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j})$, it is optimal to adjust the capacity up to the level $\underline{\mu}_j(\boldsymbol{\pi_j})$.

For an arbitrary large $\hat{\mu}_{j-1}$, we have $f(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) + \gamma_a\hat{\mu} + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]$ and the corresponding Lovász extension by $f^L(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$. Then we define $\overline{\mu}_j(\boldsymbol{\pi_j})$ such that

$$\overline{\mu}_j(\boldsymbol{\pi_j}) = \arg\max_{\hat{\mu}\in\mathbb{R}_+} \left\{f^L(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})\right\} \tag{EC.63}$$

As f^L is the piecewise linear concave envelope for f, we have $\overline{\mu}_j(\boldsymbol{\pi_j}) \in \mathcal{K}$. It is optimal for the firm to disinvest its capacity to $\overline{\mu}_j(\boldsymbol{\pi_j})$ for all $\hat{\mu}_{j-1} > \overline{\mu}_j(\boldsymbol{\pi_j})$. As $c_a \ge 0$ and $c_a \ge -\gamma_a$, it follows that $\overline{\mu}_j(\boldsymbol{\pi_j}) \ge \underline{\mu}_j(\boldsymbol{\pi_j})$. Following the concavity of the Lovász extension of the value-to-go function, it is optimal for the firm to stay put when $\underline{\mu}_j(\boldsymbol{\pi_j}) \le \hat{\mu}_{j-1} \le \overline{\mu}_j(\boldsymbol{\pi_j})$. Therefore, we have proved the optimal policy is a control band policy. □

When the capacity type is continuous, we have the value-to-go functions as follows. For all $j \in \{1, 2, ..., J-l\}$,

$$\begin{aligned} V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) &= \max_{\hat{\mu}\in\mathbb{R}_+} E\left[H_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) + V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right] \\ &= \max_{\hat{\mu}\in\mathbb{R}_+} \left\{h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - \hat{C}(\hat{\mu}_{j-1}, \hat{\mu}) + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]\right\}; \\ V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) &= 0 \text{ for } j > J-l. \end{aligned} \tag{EC.64}$$

The optimal policy for the continuous case extends the result in Proposition 4 and is shown as follows.

Proposition EC.2 **(Optimal policy for multiple adjustment with continuous capacity).** *Suppose the firm has information vector $\boldsymbol{\pi_j}$ and capacity position $\hat{\mu}_{j-1}$ at the beginning of period j. Then, the optimal capacity position, denoted by $\hat{\mu}^*(\boldsymbol{\pi_j})$, is characterized by two thresholds $\underline{\mu}_j(\boldsymbol{\pi_j})$ and $\bar{\mu}_j(\boldsymbol{\pi_j})$, such that:*

(i) If $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j})$, it is optimal for the firm to adjust the capacity position up to $\hat{\mu}^*(\boldsymbol{\pi_j}) = \underline{\mu}_j(\boldsymbol{\pi_j})$.

(ii) If $\underline{\mu}_j(\boldsymbol{\pi_j}) \le \hat{\mu}_{j-1} \le \bar{\mu}_j(\boldsymbol{\pi_j})$, it is optimal for the firm to stay put, i.e., $\hat{\mu}^*(\boldsymbol{\pi_j}) = \hat{\mu}_{j-1}$.

(iii) If $\hat{\mu}_{j-1} > \bar{\mu}_j(\boldsymbol{\pi_j})$, it is optimal for the firm to adjust the capacity position down to $\hat{\mu}^*(\boldsymbol{\pi_j}) = \bar{\mu}_j(\boldsymbol{\pi_j})$.

Proof of Proposition EC.2. The proof follows in two steps: (1) show $V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1})$ is concave in the capacity position $\hat{\mu}_{j-1}$ for all $j \le J-l+1$; (2) show the optimal policy follows the control band structure and find the lower and upper thresholds.

We first show the concavity by induction. For $j = J-l+1$, we have $V_{J-l+1}^m(\boldsymbol{\pi_{J-l+1}}, \hat{\mu}_{J-l}) = 0$, and therefore is concave in $\hat{\mu}_{J-l}$. For $j+1$, assume $V_{j+1}^m(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j)$ is concave in $\hat{\mu}_j$. It follows that $E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]$ is concave in $\hat{\mu}$ as the positive combination of concave functions is concave. Therefore, we have $h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a(\hat{\mu} - \hat{\mu}_{j-1})^+ - \gamma_a(\hat{\mu}_{j-1} - \hat{\mu})^+ + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]$ is jointly concave in $(\mu, \hat{\mu})$. For a jointly-concave function $f(x,y)$ and a convex set $\mathcal{Y}$, we have $g(x) = \max_{y\in\mathcal{Y}} f(x,y)$ is concave in x. Then it follows that

$$V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = \max_{\hat{\mu}\in\mathbb{R}_+} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a(\hat{\mu} - \hat{\mu}_{j-1})^+ - \gamma_a(\hat{\mu}_{j-1} - \hat{\mu})^+ + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right] \right\}$$

is concave in $\hat{\mu}_{j-1}$.

We next show the optimal policy follows the control band structure. For $\hat{\mu}_{j-1} = 0$, we define

$$\underline{\mu}_j(\boldsymbol{\pi_j}) \triangleq \underset{\hat{\mu}\in\mathbb{R}_+}{\arg\max} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a\hat{\mu} + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right] \right\} \tag{EC.65}$$

For any $\hat{\mu}_{j-1}$ that is sufficiently small, i.e., $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j})$, it is optimal to adjust the capacity up to the level $\underline{\mu}_j(\boldsymbol{\pi_j})$.

For an arbitrary large $\hat{\mu}_{j-1}$, we define

$$\overline{\mu}_j(\boldsymbol{\pi_j}) \triangleq \underset{\hat{\mu}\in\mathbb{R}_+}{\arg\max} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) + \gamma_a\hat{\mu} + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right] \right\} \tag{EC.66}$$

It is optimal for the firm to disinvest its capacity to $\overline{\mu}_j(\boldsymbol{\pi_j})$ for all $\hat{\mu}_{j-1} > \overline{\mu}_j(\boldsymbol{\pi_j})$. As $c_a \ge 0$ and $c_a \ge -\gamma_a$, it follows that $\overline{\mu}_j(\boldsymbol{\pi_j}) \ge \underline{\mu}_j(\boldsymbol{\pi_j})$. Following the concavity of the value-to-go function, it is optimal for the firm to stay put when $\underline{\mu}_j(\boldsymbol{\pi_j}) \le \mu \le \overline{\mu}_j(\boldsymbol{\pi_j})$. Therefore, we have proved the optimal policy is a control band policy. □

Proof of Proposition 5. We first note that the likelihood ratio order implies first order stochastic dominance. To simplify the notations, we define the two functions

$$\begin{aligned} G_j^a(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) &= h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a(\hat{\mu} - \hat{\mu}_{j-1})^+ - \gamma_a(\hat{\mu}_{j-1} - \hat{\mu})^+ + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right] \\ G_j^s(\boldsymbol{\pi_j}, \hat{\mu}) &= h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]. \end{aligned} \tag{EC.67}$$

We observe that $V_j^m(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = \max_{\hat{\mu}} G_j^a(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$. We also have the fact that $G_j^a(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$ is concave in $\hat{\mu}$ from Proposition EC.2. To show the two thresholds increase in the information vector $\boldsymbol{\pi_j}$, it is sufficient to show that for $j = 1, ..., J-l$, $\frac{\partial G_j^a}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$ increases in $\boldsymbol{\pi_j}$ for any $\hat{\mu} \ne \hat{\mu}_{j-1}$, and $\frac{\partial G_j^s}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu})$, which is a special case when $\hat{\mu} = \hat{\mu}_{j-1}$, increases in $\boldsymbol{\pi_j}$. We prove this by induction. We present the result for $G_j^a(\cdot)$ as the proof for $G_j^s(\cdot)$ is identical.

To establish induction basis, let $j = J - l$. If $\hat{\mu} \le \hat{\mu}_{J-l-1}$, we have

$$\frac{\partial G^a_{J-l}}{\partial \hat{\mu}}(\boldsymbol{\pi_{J-l}}, \hat{\mu}_{J-l-1}, \hat{\mu}) = \sum_{i=1}^{I} \pi_{J-l,i}[c_1\tau(1 - F_J(\hat{\mu}\tau|\theta_i)) - c_0\tau] + \gamma_a \tag{EC.68}$$

Therefore, as $D_j|\theta_i \preceq_{st} D_j|\theta_{\hat{i}}$ for $i \le \hat{i}$, we have $\frac{\partial G^a_{J-l}}{\partial \hat{\mu}}(\boldsymbol{\pi_{J-l}}, \hat{\mu}_{J-l-1}, \hat{\mu})$ increases in $\boldsymbol{\pi_{J-l}}$ when $\hat{\mu} < \hat{\mu}_{J-l-1}$. A similar argument establishes the result for the case $\hat{\mu} > \hat{\mu}_{J-l-1}$.

Suppose that the result hold for all $t = j+1, ..., J-l$. Thus, at period $j+1$, $\frac{\partial G^a_{j+1}}{\partial \hat{\mu}}(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j, \hat{\mu})$ increases in $\boldsymbol{\pi_{j+1}}$ for all $\hat{\mu} \ne \hat{\mu}_j$. From the induction hypothesis, the two switching curves $\underline{\mu}_{j+1}(\boldsymbol{\pi_{j+1}})$ and $\bar{\mu}_{j+1}(\boldsymbol{\pi_{j+1}})$ increase in $\boldsymbol{\pi_{j+1}}$ as well. We now show that $\frac{\partial G^a_j}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$ increases in $\boldsymbol{\pi_j}$ for $\hat{\mu} \ne \hat{\mu}_{j-1}$.

For any $\hat{\mu} < \hat{\mu}_{j-1}$, we have

$$\frac{\partial G^a_j}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu}) = \sum_{i=1}^{I} \pi_{j,i}[c_1\tau(1 - F_{j+l}(\hat{\mu}\tau|\theta_i)) - c_0\tau] + \gamma_a + E\left[\frac{\partial V^m_{j+1}}{\partial \hat{\mu}}(\boldsymbol{\Pi_{j+1}}, \hat{\mu})\middle| \boldsymbol{\pi_j}\right] \tag{EC.69}$$

The expression for $\frac{\partial G^a_j}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$ when $\hat{\mu} > \hat{\mu}_{j-1}$ is similar. In order to show equation (EC.69) increases in $\boldsymbol{\pi_j}$, we need to show that $E\left[\frac{\partial V^m_{j+1}}{\partial \hat{\mu}}(\boldsymbol{\Pi_{j+1}}, \hat{\mu})\middle| \boldsymbol{\pi_j}\right]$ increases in $\boldsymbol{\pi_j}$, which is shown in two steps.

We first observe that

$$\frac{\partial V^m_{j+1}}{\partial \hat{\mu}_j}(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j) = \begin{cases} c_a & \text{if } \hat{\mu}_j < \underline{\mu}_{j+1}(\boldsymbol{\pi_{j+1}}) \\ \frac{\partial G^s_{j+1}}{\partial \hat{\mu}}(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j) & \text{if } \underline{\mu}_{j+1}(\boldsymbol{\pi_{j+1}}) \le \hat{\mu}_j \le \bar{\mu}_{j+1}(\boldsymbol{\pi_{j+1}}) \\ -\gamma_a & \text{if } \hat{\mu}_j > \bar{\mu}_{j+1}(\boldsymbol{\pi_{j+1}}) \end{cases} \tag{EC.70}$$

Notice that $\frac{\partial V^m_{j+1}}{\partial \hat{\mu}_j}(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j)$ is continuously decreasing in $\hat{\mu}_j$, and $\underline{\mu}_{j+1}(\boldsymbol{\pi_{j+1}})$ and $\bar{\mu}_{j+1}(\boldsymbol{\pi_{j+1}})$ increase in $\boldsymbol{\pi_{j+1}}$. From the induction hypothesis, $\frac{\partial V^m_{j+1}}{\partial \hat{\mu}_j}(\boldsymbol{\pi_{j+1}}, \hat{\mu}_j)$ increases in $\boldsymbol{\pi_{j+1}}$.

Next, we show that $\boldsymbol{\Pi_{j+1}}|\boldsymbol{\pi_j}$ increases in $\boldsymbol{\pi_j}$ in the first order stochastic dominance sense. Notice that as $\boldsymbol{\pi_j}$ increases in the likelihood ratio order, $\boldsymbol{\pi_{j+1}}$ increases in the likelihood ratio order, which implies that $\boldsymbol{\pi_{j+1}}$ increases in the first order stochastic dominance. In addition, we have the stochastic dominance relationship among the demand types, i.e., $D_j|\theta_i \preceq_{st} D_j|\theta_{\hat{i}}$ for $i \le \hat{i}$. Therefore, $\boldsymbol{\Pi_{j+1}}|\boldsymbol{\pi_j}$ stochastically increases in $\boldsymbol{\pi_j}$. Thus, the fact that $E\left[\frac{\partial V^m_{j+1}}{\partial \hat{\mu}}(\boldsymbol{\Pi_{j+1}}, \hat{\mu})\middle| \boldsymbol{\pi_j}\right]$ increases in $\boldsymbol{\pi_j}$ immediately follows. It then follows that $\frac{\partial G^a_j}{\partial \hat{\mu}}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}, \hat{\mu})$ increases in $\boldsymbol{\pi_j}$ for $\hat{\mu} < \hat{\mu}_{j-1}$. A similar argument proves the case for $\hat{\mu} > \hat{\mu}_{j-1}$. Therefore, the result holds for period j. □

Proof of Proposition 6. Similar to the proof of Proposition 3, we first find an upper bound of the regret by finding a lower bound of $V^{ms}_{0,n}$. To simplify the notations, we define the firm's expected profits under the multi-step heuristic (given the demand type λ_i) in different periods as follows. We still use $\hat{\lambda}_{i,j\tau_n}$ to denote $\frac{D_j|\theta_{i,n}}{n\tau_n}$. First, in period $l_n + 1$, the firm's capacity is still the initial capacity μ_0, and we have

$$W_{s,n}(\lambda_i) \triangleq E\left\{pn\hat{\lambda}_{i,(l_n+1)\tau_n}\tau_n - c_1 n\left(\hat{\lambda}_{i,(l_n+1)\tau_n} - \mu_0\right)^+ \tau_n - c_0 n\mu_0\tau_n\right\} \tag{EC.71}$$

Second, during period l_n+2 and $l_n+2^{K_n}-1$, the firm's capacity level is updated according to the heuristic. Observing that $E[\bar{\lambda}_{i,\kappa}]=\lambda_i$, we have the firm's expected profits as

$$\begin{aligned} W_{m,n}(\lambda_i) \triangleq & E\sum_{\kappa=1}^{K_n-1}\left\{\sum_{j=l_n+2^\kappa}^{l_n+2^{\kappa+1}-1}\left[pn\hat{\lambda}_{i,j\tau_n}\tau_n - c_1 n\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,\kappa}\right)^+\tau_n - c_0 n\bar{\lambda}_{i,\kappa}\tau_n\right] - \hat{C}\left(n\bar{\lambda}_{i,\kappa-1}, n\bar{\lambda}_{i,\kappa}\right)\right\} \\ =& \sum_{\kappa=1}^{K_n-1}\left\{(p-c_0)n\lambda_i 2^\kappa\tau_n - c_1 n\tau_n\sum_{j=l_n+2^\kappa}^{l_n+2^{\kappa+1}-1} E\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,\kappa}\right)^+ - E\left[\hat{C}\left(n\bar{\lambda}_{i,\kappa-1}, n\bar{\lambda}_{i,\kappa}\right)\right]\right\} \end{aligned} \tag{EC.72}$$

Finally, during period $l_n+2^{K_n}$ and J_n, the firm makes the last adjustment, and the capacity maintains at this level for the rest of the time horizon. Then we have

$$\begin{aligned} W_{l,n}(\lambda_i) \triangleq & E\left\{\sum_{j=l_n+2^{K_n}}^{J_n}\left[pn\hat{\lambda}_{i,j\tau_n}\tau_n - c_1 n\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,K_n}\right)^+\tau_n - c_0 n\bar{\lambda}_{i,K_n}\tau_n\right] - \hat{C}\left(n\bar{\lambda}_{i,K_n-1}, n\bar{\lambda}_{i,K_n}\right)\right\} \\ =&(p-c_0)n\lambda_i\left(T-l_t-\left(2^{K_n}-1\right)\tau_n\right) - c_1 n\tau_n\sum_{j=l_n+2^{K_n}}^{J_n} E\Big(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,K_n}\Big)^+ \\ & - E\left[\hat{C}\left(n\bar{\lambda}_{i,K_n-1}, n\bar{\lambda}_{i,K_n}\right)\right] \end{aligned} \tag{EC.73}$$

We show analysis for the case where $\lambda_i \geq \mu_0$ for all i, as the analysis for the other case is similar. Because $W_{s,n}(\lambda_i)\geq 0$ as $p\geq c_1>c_0$, we have

$$V_{0,n}^{ms} = \sum_{i=1}^{I}\pi_{1,i}\Big\{W_{s,n}(\lambda_i)+W_{m,n}(\lambda_i)+W_{l,n}(\lambda_i)\Big\} \geq \sum_{i=1}^{I}\pi_{1,i}\Big\{W_{m,n}(\lambda_i)+W_{l,n}(\lambda_i)\Big\} \tag{EC.74}$$

Therefore, by equation (20), we have an upper bound of the regret as follows

$$\begin{aligned} R_n^{\pi_n} &= 1 - V_{0,n}^{ms}/V_{0,n}^d \\ &\leq \frac{1}{V_{0,n}^d}\sum_{i=1}^{I}\pi_{1,i}\left\{(p-c_0)n\lambda_i(T-l_t) - \hat{C}(n\mu_0, n\lambda_i) - W_{m,n}(\lambda_i) - W_{l,n}(\lambda_i)\right\} \\ &= \frac{1}{V_{0,n}^d}\sum_{i=1}^{I}\pi_{1,i}\left\{\begin{array}{l}(p-c_0)n\lambda_i\tau_n - \hat{C}(n\mu_0, n\lambda_i) + \sum_{\kappa=1}^{K_n} E\left[\hat{C}\left(n\bar{\lambda}_{i,\kappa-1}, n\bar{\lambda}_{i,\kappa}\right)\right] \\ +c_1 n\tau_n\left[\sum_{\kappa=1}^{K_n-1}\sum_{j=l_n+2^\kappa}^{l_n+2^{\kappa+1}-1} E\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,\kappa}\right)^+ + \sum_{j=l_n+2^{K_n}}^{J_n} E\Big(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,K_n}\Big)^+\right]\end{array}\right\} \end{aligned} \tag{EC.75}$$

To find an upper bound for the right hand side of equation (EC.75), we need to find an upper bound for $E\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,\kappa}\right)^+$, $E(\bar{\lambda}_{i,\kappa}-\bar{\lambda}_{i,\kappa-1})^+$, and $E(\bar{\lambda}_{i,\kappa-1}-\bar{\lambda}_{i,\kappa})^+$ respectively. Note that $E[\bar{\lambda}_{i,\kappa}]=\lambda_i$ and $Var[\bar{\lambda}_{i,\kappa}]=\frac{\sigma^2\lambda_i}{n(2^\kappa-1)\tau_n}$. We use C_i to represent a constant which is independent of n and τ_n for all i. By Proposition EC.1 and the inequality of (EC.33), we have

$$E\left(\hat{\lambda}_{i,j\tau_n}-\bar{\lambda}_{i,\kappa}\right)^+ \leq E\left(\hat{\lambda}_{i,j\tau_n}-\lambda_i\right)^+ + E\left(\lambda_i-\bar{\lambda}_{i,\kappa}\right)^+ \leq \frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n\tau_n}} + \frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n(2^\kappa-1)\tau_n}} \leq \frac{C_3}{\sqrt{n\tau_n}} \tag{EC.76}$$

For $\kappa=1$, we have

$$E\left(\bar{\lambda}_{i,1}-\bar{\lambda}_{i,0}\right)^{+}=E\left(\bar{\lambda}_{i,1}-\mu_0\right)^{+}\le E\left(\bar{\lambda}_{i,1}-\lambda_i\right)^{+}+(\lambda_i-\mu_0)^{+}\le\frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n\tau_n}}+(\lambda_i-\mu_0)^{+} \tag{EC.77}$$

$$E\left(\bar{\lambda}_{i,0}-\bar{\lambda}_{i,1}\right)^{+}=E\left(\mu_0-\bar{\lambda}_{i,1}\right)^{+}\le(\mu_0-\lambda_i)^{+}+E\left(\lambda_i-\bar{\lambda}_{i,1}\right)^{+}\le(\mu_0-\lambda_i)^{+}+\frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n\tau_n}} \tag{EC.78}$$

By equation (18), for $\kappa\ge 2$, we have

$$\begin{aligned}
E\left(\bar{\lambda}_{i,\kappa}-\bar{\lambda}_{i,\kappa-1}\right)^{+}&=E\left(\frac{\bar{\lambda}_{i,\kappa-1}n(2^{\kappa-1}-1)\tau_n+N(n\lambda_i(2^{\kappa}-1)\tau_n)-N(n\lambda_i(2^{\kappa-1}-1)\tau_n)}{n(2^{\kappa}-1)\tau_n}-\bar{\lambda}_{i,\kappa-1}\right)^{+}\\
&=E\left(\frac{N(n\lambda_i(2^{\kappa}-1)\tau_n)-N(n\lambda_i(2^{\kappa-1}-1)\tau_n)-\bar{\lambda}_{i,\kappa-1}n2^{\kappa-1}\tau_n}{n(2^{\kappa}-1)\tau_n}\right)^{+}\\
&\le E\left(\frac{N(n\lambda_i(2^{\kappa}-1)\tau_n)-N(n\lambda_i(2^{\kappa-1}-1)\tau_n)}{n2^{\kappa-1}\tau_n}-\lambda_i\right)^{+}+E\left(\lambda_i-\bar{\lambda}_{i,\kappa-1}\right)^{+}\\
&\le\frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n2^{\kappa-1}\tau_n}}+\frac{\sigma\sqrt{\lambda_i}}{2\sqrt{n(2^{\kappa-1}-1)\tau_n}}\le\frac{C_5}{\sqrt{n2^{\kappa}\tau_n}} \qquad \text{(EC.79)}\\
E\left(\bar{\lambda}_{i,\kappa-1}-\bar{\lambda}_{i,\kappa}\right)^{+}&\le\frac{C_6}{\sqrt{n2^{\kappa}\tau_n}} \qquad \text{(EC.80)}
\end{aligned}$$

We next apply the inequality (EC.76) to (EC.80) to the right hand side of equation (EC.75), and gather the items by the outsourcing costs and capacity adjustment costs respectively. Then we obtain an upper bound of the regret as follows

$$\begin{aligned}
\text{RHS of (EC.75)}&\le C_7\tau_n+C_8\frac{c_1}{\sqrt{n\tau_n}}+C_9\sum_{\kappa=1}^{K_n}\frac{\max(c_a,\gamma_a^{+})}{\sqrt{n2^{\kappa}\tau_n}}\\
&\le C_7\tau_n+\frac{C_{10}}{\sqrt{n\tau_n}} \qquad \text{(EC.81)}
\end{aligned}$$

The last inequality follows the fact that K_n satisfies that $(2^{K_n+1}-1)\tau_n\le T-l_t$.

By setting $\tau_n\asymp n^{-\frac{1}{3}}$, we obtain an upper bound of the regret on the order of $n^{-\frac{1}{3}}$. □

Remark: The exponentially increasing time between two consecutive adjustments is important in establishing the upper bound in the order of $n^{-1/3}$. To illustrate this, we alternatively consider another heuristic, where the time between two consecutive adjustments is fixed as $\hat{n}\tau_n$, $\hat{n}\in\mathbb{N}^{+}$. We denote the regret under this heuristic as R_n^{fa}. Following the same logic as in the proof of Proposition 6, it can be obtained that the upper bound of the regret satisfies the following

$$\begin{aligned}
R_n^{fa}&\le C_7\tau_n+C_{11}\frac{c_1}{\sqrt{n\tau_n}}+C_{12}\frac{\max(c_a,\gamma_a^{+})}{\sqrt{n\hat{n}\tau_n}}\frac{T}{\hat{n}\tau_n}\\
&\le C_7\tau_n+C_{13}\frac{1}{\sqrt{n\tau_n^3}} \qquad \text{(EC.82)}
\end{aligned}$$

In this case, the firm should set $\tau_n\asymp n^{-1/5}$ and yield an upper bound in the order of $n^{-1/5}$. It cannot tighten the upper bound to the order of $n^{-1/3}$, because the capacity adjustment is too frequent and the adjustment cost is too high.

EC.2.2.2. Remark

Our results can be generalized to the case with fixed capacity cost, the capacity reduction is limited by the current available capacity, and different leadtimes to increase and decreases capacity as follows.

Fixed cost. If there exist fixed costs with capacity change, we can show that the optimal policy has state-dependent regions where the firm can alternate between increasing and staying put (or between decreasing and staying put) multiple times as the initial capacity increases: e.g., ISDSD and ISISD. This shows that, with fixed costs, even the fundamental structure such as "the optimal policy is of an ISD type" no longer holds. This is because concavity (its preservation under maximum over a convex set), which is the original machinery of proving the optimality of an ISD policy, no longer holds with fixed costs. To characterize the structure of the optimal policy, we use weak (K_1, K_2)-concavity, developed in Semple (2007). For a given information vector, the structure of the optimal policy is now characterized with several thresholds with respect to the initial capacity. Specifically, as the initial capacity increases, the optimal policy may switch between investing and staying put before disinvestment region starts. Likewise, the optimal policy may alternate between staying put and disinvestment multiple times. We show that these thresholds are dependent on the information vector, which represents the firm's belief about demand types. Thus, we generalize the result of Semple (2007) to the case with an unknown demand type.

The weak (K_1, K_2)-concavity developed in Semple (2007) is as follows, where K_1 and K_2 are given nonnegative constants.

Definition EC.2 (Semple (2007): Weak (K_1, K_2)-concavity). A continuous function $f(x)$ is weakly (K_1, K_2)-concave on the interval $[0, U]$ if and only if for any two points $x, y \in [0, U]$ with $x < y$ and any $\lambda \in [0, 1]$,

$$f((1-\lambda)x + \lambda y) \geq (1-\lambda)(f(x) - K_2) + \lambda(f(y) - K_1) \tag{EC.83}$$

The concept of weak (K_1, K_2)-concavity allows us to characterize the structure of optimal policy. If the value-to-go function in one period satisfies the weak (K_1, K_2)-concavity, the structure of the optimal policy can be shown as the one in Proposition EC.3 below. In addition, it can be shown that the weak (K_1, K_2)-concavity preserves in the dynamic program, and therefore, we are able to completely characterize the optimal policy with respect to initial capacity. Below we describe the details of the optimal policy.

We first define the value-to-go function. We define the cost to adjust capacity from μ to μ' with fixed costs as $\check{C}(\mu, \mu') \triangleq c_a(\mu' - \mu)^+ + \overline{K}\mathbb{1}_{\{\mu' - \mu > 0\}} + \gamma_a(\mu - \mu')^+ + \underline{K}\mathbb{1}_{\{\mu - \mu' > 0\}}$. Let the superscript

mf denote the case of multiple capacity adjustment with a fixed cost and U denote a large constant, and we have

$$\begin{aligned}
&V_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) = \max_{\hat{\mu}\in[0,U]}\left\{h_{j+l}\left(\boldsymbol{\pi_j},\hat{\mu}\right) - \check{C}\left(\hat{\mu}_{j-1},\hat{\mu}\right) + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}},\hat{\mu})|\boldsymbol{\pi_j}\right]\right\};\\
&V_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) = 0 \text{ for } j > J-l.
\end{aligned} \tag{EC.84}$$

To simplify notations, we define the value-to-go function assuming the firm can only increase its capacity, $\overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1})$, and the value-to-go function assuming the firm can only decrease its capacity, $\underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1})$ in period j as follows.

$$\begin{aligned}
\overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) &= \max_{\hat{\mu}\in[\hat{\mu}_{j-1},U]}\left\{h_{j+l}\left(\boldsymbol{\pi_j},\hat{\mu}\right) - c_a(\hat{\mu}-\hat{\mu}_{j-1}) - \overline{K}\mathbb{1}_{\{\hat{\mu}-\hat{\mu}_{j-1}>0\}} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}},\hat{\mu})|\boldsymbol{\pi_j}\right]\right\}\\
\underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) &= \max_{\hat{\mu}\in[0,\hat{\mu}_{j-1}]}\left\{h_{j+l}\left(\boldsymbol{\pi_j},\hat{\mu}\right) - \gamma_a(\hat{\mu}_{j-1}-\hat{\mu}) - \underline{K}\mathbb{1}_{\{\hat{\mu}_{j-1}-\hat{\mu}>0\}} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}},\hat{\mu})|\boldsymbol{\pi_j}\right]\right\}
\end{aligned} \tag{EC.85}$$

It is immediate that $V_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) = \max\{\overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}), \underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1})\}$.

To simplify description of the optimal policy, we define the following critical values in period j.

$$\begin{aligned}
&B_j(\boldsymbol{\pi_j}) \triangleq \sup\{\hat{\mu}\in\arg\max_{\hat{\mu}\in[0,U]}\left[h_{j+l}\left(\boldsymbol{\pi_j},\hat{\mu}\right) - c_a\hat{\mu} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}},\hat{\mu})|\boldsymbol{\pi_j}\right]\right]\},\\
&S_j(\boldsymbol{\pi_j}) \triangleq \inf\{\hat{\mu}\in\arg\max_{\hat{\mu}\in[0,U]}\left[h_{j+l}\left(\boldsymbol{\pi_j},\hat{\mu}\right) + \gamma_a\hat{\mu} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}},\hat{\mu})|\boldsymbol{\pi_j}\right]\right]\},\\
&b_j(\boldsymbol{\pi_j}) \triangleq \sup\{\mu : \overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) > \underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}), \forall\hat{\mu}_{j-1}\in[0,\mu)\}, \text{ if } \{\} \text{ empty, set } b_j(\boldsymbol{\pi_j}) = 0,\\
&b_j^*(\boldsymbol{\pi_j}) \triangleq \inf\{\mu : \overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) \le \underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}), \forall\hat{\mu}_{j-1}\in[\mu,U]\},\\
&s_j^*(\boldsymbol{\pi_j}) \triangleq \sup\{\mu : \overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) \ge \underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}), \forall\hat{\mu}_{j-1}\in[0,\mu]\},\\
&s_j(\boldsymbol{\pi_j}) \triangleq \inf\{\mu : \overline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}) < \underline{V}_j^{mf}(\boldsymbol{\pi_j},\hat{\mu}_{j-1}), \forall\hat{\mu}_{j-1}\in(\mu,U]\}, \text{ if } \{\} \text{ empty, set } s_j(\boldsymbol{\pi_j}) = U.
\end{aligned} \tag{EC.86}$$

That is, given the information vector $\boldsymbol{\pi_j}$, $B_j(\boldsymbol{\pi_j})$ and $S_j(\boldsymbol{\pi_j})$ represent the largest global invest-up-to capacity position and the smallest global disinvest-down-to capacity position respectively. With respect to the optimal policy in period j, we have $b_j(\boldsymbol{\pi_j})$ is the largest value below which the firm always invests, $b_j^*(\boldsymbol{\pi_j})$ is the smallest value above which the firm never invests, $s_j^*(\boldsymbol{\pi_j})$ is the largest value below which the firm never disinvests, and $s_j(\boldsymbol{\pi_j})$ is the smallest value above which the firm always disinvests. Following a similar proof of Theorem 2, 3, 4, and 5 in Semple (2007), we establish the following structure of the optimal policy.

Proposition EC.3 **(Optimal policy for multiple adjustment with fixed cost)**. *In period j, given information vector $\boldsymbol{\pi_j}$ and capacity position $\hat{\mu}_{j-1}$, there are two cases.*

Case 1: $\overline{K} \le \underline{K}$ (disinvestment has a higher fixed cost.)

(i) $\hat{\mu}_{j-1}\in[0, b_j(\boldsymbol{\pi_j}))$, the firm invests up to $B_j(\boldsymbol{\pi_j})$.

(ii) $\hat{\mu}_{j-1} \in [b_j(\boldsymbol{\pi_j}), b_j^*(\boldsymbol{\pi_j}))$, *the firm stays put if* $\overline{V}_j^{mf}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) + E[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1})|\boldsymbol{\pi_j}]$. *Otherwise the firm invests up to* $\tilde{B}_j(\boldsymbol{\pi_j}) \in \arg\max_{\hat{\mu}\in[\hat{\mu}_{j-1},U]} \left\{h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a(\hat{\mu} - \hat{\mu}_{j-1}) - \overline{K}\mathbb{1}_{\{\hat{\mu}-\hat{\mu}_{j-1}>0\}} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]\right\}$.

(iii) $\hat{\mu}_{j-1} \in [b_j^*(\boldsymbol{\pi_j}), s_j^*(\boldsymbol{\pi_j})]$, *the firm stays put.*

(iv) $\hat{\mu}_{j-1} \in [s_j^*(\boldsymbol{\pi_j}), s_j(\boldsymbol{\pi_j})]$, *the firm stays put if* $\underline{V}_j^{mf}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) + E[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1})|\boldsymbol{\pi_j}]$. *Otherwise the firm disinvests down to* $S_j(\boldsymbol{\pi_j})$.

(v) $\hat{\mu}_{j-1} \in [s_j(\boldsymbol{\pi_j}), U]$, *the firm should disinvests to* $S_j(\boldsymbol{\pi_j})$.

Case 2: $\overline{K} > \underline{K}$ *(investment has a higher fixed cost.)*

(i) $\hat{\mu}_{j-1} \in [0, b_j(\boldsymbol{\pi_j}))$, *the firm invests up to* $B_j(\boldsymbol{\pi_j})$.

(ii) $\hat{\mu}_{j-1} \in [b_j(\boldsymbol{\pi_j}), b_j^*(\boldsymbol{\pi_j}))$, *the firm stays put if* $\overline{V}_j^{mf}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) + E[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1})|\boldsymbol{\pi_j}]$. *Otherwise the firm invests up to* $B_j(\boldsymbol{\pi_j})$.

(iii) $\hat{\mu}_{j-1} \in [b_j^*(\boldsymbol{\pi_j}), s_j^*(\boldsymbol{\pi_j})]$, *the firm stays put.*

(iv) $\hat{\mu}_{j-1} \in [s_j^*(\boldsymbol{\pi_j}), s_j(\boldsymbol{\pi_j})]$, *the firm stays put if* $\underline{V}_j^{mf}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) = h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}_{j-1}) + E[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu}_{j-1})|\boldsymbol{\pi_j}]$. *Otherwise the firm disinvests down to* $\tilde{S}_j(\boldsymbol{\pi_j}) \in \arg\max_{\hat{\mu}\in[\hat{\mu}_{j-1},U]} \left\{h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - \gamma_a(\hat{\mu}_{j-1} - \hat{\mu}) - \underline{K}\mathbb{1}_{\{\hat{\mu}_{j-1}-\hat{\mu}>0\}} + E\left[V_{j+1}^{mf}(\boldsymbol{\Pi_{j+1}}, \hat{\mu})|\boldsymbol{\pi_j}\right]\right\}$.

(v) $\hat{\mu}_{j-1} \in [s_j(\boldsymbol{\pi_j}), U]$, *the firm should disinvests to* $S_j(\boldsymbol{\pi_j})$.

Note that we have established the optimal policy for the continuous capacity case. When the capacity is discrete, it remains an open question about how to extend the weak (K_1, K_2)-concavity to the discrete case and characterize the optimal policy.

Leadtime. When the capacity reduction is limited by the current available capacity, we show that the optimal policy is still of the ISD type, where the thresholds will depend on the pipeline capacity to be installed. The details are provided in Section EC.2.3.

In the baseline model, we assume the leadtime to adjust the capacity is l periods, which remains the same for capacity investment and disinvestment. Similar to the discussion in Section EC.2.1.2, one may expect that the leadtime to invest in capacity, $\bar{l}$, might be different from the leadtime to disinvest, $\underline{l}$. In what follows, we establish that when the firm does not actively update its belief about demand, i.e., the firm's belief about demand types remains as $\boldsymbol{\pi_1}$, the decision problem with different leadtimes can be recast as an equivalent one with the same leadtime following a similar discussion in Ye and Duenyas (2007). We consider the case where $\bar{l} > \underline{l}$, and the analysis for the other case is similar.

Let μ_j denote the capacity level in period j after the capacity invested in period $j - \bar{l}$ is installed and the capacity disinvested in period $j - \underline{l}$ is salvaged. Also define c_a^j be the cost to add a unit of capacity in period $j + \bar{l}$ and γ_a^j be the cost to disinvest one unit of capacity in period $j + \underline{l}$. Then the firm's capacity adjustment decision problem is to find a sequence of $(\mu_1,...,\mu_J)$ which maximizes the total profit over the horizon:

$$\max_{(\mu_1,\dots,\mu_J)} \sum_{j=1}^{J} E\left[h_j(\boldsymbol{\pi_1},\mu_j) - c_a^{j-\bar{l}}(\mu_j-\mu_{j-1})^+ - \gamma_a^{j-\underline{l}}(\mu_{j-1}-\mu_j)^+\right] \tag{EC.87}$$

where we have $c_a^{j-\bar{l}} = M$ for all $j=1,2,...,\bar{l}$ and $c_a^{j-\bar{l}} = c_a$ otherwise; $\gamma_a^{j-\underline{l}} = M$ for all $j=1,2,...,\underline{l}$ and $\gamma_a^{j-\underline{l}} = \gamma_a$ otherwise. Here M stands for a very large number. It is immediate that $\mu_j = \mu_0$ for $j=1,2,...,\underline{l}$, $\mu_j \le \mu_{j-1}$ for all $j=\underline{l}+1,...,\bar{l}$, and $\mu_j \in \mathcal{K}$ for $j=\bar{l}+1,...,J$.

Then we define the equivalent optimal value-to-go function as follows: for all $j=1,2,...,J$,

$$V_j^m(\boldsymbol{\pi_1},\check{\mu}_{j-1}) = \max_{\check{\mu}\in\mathcal{K}}\left\{h_j(\boldsymbol{\pi_1},\check{\mu}) - c_a^{j-\bar{l}}(\check{\mu}-\check{\mu}_{j-1})^+ - \gamma_a^{j-\underline{l}}(\check{\mu}_{j-1}-\check{\mu})^+ + E\left[V_{j+1}^m(\boldsymbol{\pi_1},\check{\mu})\right]\right\};$$
$$V_{J+1}^m(\boldsymbol{\pi_1},\check{\mu}_J) = 0. \tag{EC.88}$$

Note the dynamic program has a similar structure as the one in Section 5, we can obtain a similar structure of the optimal policy following the analysis of Proposition 4 and EC.2. With the optimal solution of $(\check{\mu}_1^*,...,\check{\mu}_J^*)$, the decision the firm should make in period j is as follows. For $j=1,2,...,J-\bar{l}$, invests $(\check{\mu}_j^*-\check{\mu}_{j-1}^*)^+$; for $j=1,2,...,J-\underline{l}$, disinvests $(\hat{\mu}_{j-1}^*-\hat{\mu}_j^*)^+$.

When the firm actively updates its belief about demand, the capacity adjustment decision depends on the firm's information about demand which may be different for increasing and decreasing the capacity for a period because these decisions are made at different times. By tracking a state vector of the capacity waiting to be installed or disposed between the lead time $\bar{l}$ and $\underline{l}$ as well as the belief about demand types, it is possible to show that the optimal capacity adjustment decision still follows a state-dependent ISD policy in this case.

EC.2.3. Optimal policy when capacity reduction is limited by current capacity

In this section, we derive the firm's optimal policy to adjust its capacity, when the current available capacity is the maximum capacity available for making capacity reduction decisions. The model setting is identical to the one in Section 5, except that the firm cannot reduce its capacity beyond the current available capacity level.

As the firm can only reduce its capacity from its current available capacity level, the information states regarding capacity include the current capacity level as well as those in the pipeline. We denote the capacity vector in period j as $\boldsymbol{\mu_j} = (\mu_{j,0},\mu_{j,1},...,\mu_{j,l})$ where $\mu_{j,0}$ is the current available capacity and $\mu_{j,k}$ indicates the capacity increment to be installed or disposed in k periods. As there is a leadtime of l periods in adjusting its capacity, the capacity adjustment decision (increasing or decreasing capacity) will only affect the firm's profit after the l periods of leadtime, and the firm needs to decide the total available capacity after the leadtime, which is denoted by the capacity position, $\hat{\mu}$. For simplicity, we consider the case where the capacity is continuous. Therefore, we formulate the firm's optimal value-to-go function as follows: for all $j \in \{1,2,...,J-l\}$,

$$V_j^m(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = \max_{\hat{\mu} \geq \sum_{i=1}^{l} \mu_{j-1,k}} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - \hat{C}\left(\sum_{i=0}^{l} \mu_{j-1,k}, \hat{\mu}\right) + E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j}) | \boldsymbol{\pi_j}\right] \right\};$$
$$V_j^m(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = 0 \text{ for } j > J-l. \tag{EC.89}$$

The capacity states evolve as follows:

$$\mu_{j,0} = \mu_{j-1,0} + \mu_{j-1,1} \text{ for } k = 1,...,l-1$$
$$\mu_{j,k} = \mu_{j-1,k+1} \text{ for } k = 1,...,l-1$$
$$\mu_{j,l} = \hat{\mu} - \sum_{i=0}^{l} \mu_{j-1,k} \text{ for } k = l. \tag{EC.90}$$

The constraint $\hat{\mu} \geq \sum_{i=1}^{l} \mu_{j-1,k}$ specifies that the total available capacity l periods later should be greater than all the capacity increment waiting to be installed or disposed, which means that if the capacity is reduced in the current period, the maximum capacity reduction cannot exceed the current available capacity.

With this additional constraint, we show in the next result that, for a given information vector, the optimal policy is also of a control band type. This echoes the result in Section 5.

PROPOSITION EC.4 **(Optimal policy for limited capacity reduction)**. *Let* $\boldsymbol{\pi_j}$ *be the information vector and* $\hat{\mu}_{j-1} \triangleq \sum_{i=0}^{l} \mu_{j-1,k}$ *be capacity position at the beginning of period* j*. Then, the optimal capacity position, denoted by* $\hat{\mu}^*(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*, is characterized by two thresholds* $\underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$ *and* $\bar{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*, such that:*
(i) If $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*, it is optimal for the firm to adjust the capacity position up to* $\hat{\mu}^*(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = \underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*.*
(ii) If $\underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) \leq \hat{\mu}_{j-1} \leq \bar{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*, it is optimal for the firm to stay put, i.e.,* $\hat{\mu}^*(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = \hat{\mu}_{j-1}$*.*
(iii) If $\hat{\mu}_{j-1} > \bar{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*, it is optimal for the firm to adjust the capacity position down to* $\hat{\mu}^*(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = \bar{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$*.*

Proof of Proposition EC.4. The proof is similar to the proof of EC.2 and follows in two steps: (1) show $V_j^m(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$ is concave in the capacity vector $\boldsymbol{\mu_{j-1}}$ for all $j \leq J-l+1$; (2) show the optimal policy follows the control band structure and find the lower and upper thresholds.

We first show the concavity by induction. For $j = J-l+1$, we have $V_{J-l+1}^m(\boldsymbol{\pi_{J-l+1}}, \boldsymbol{\mu_{J-l}}) = 0$, and therefore is concave in $\boldsymbol{\mu_{J-l}}$. For $j+1$, assume $V_{j+1}^m(\boldsymbol{\pi_{j+1}}, \boldsymbol{\mu_j})$ is concave in $\boldsymbol{\mu_j}$. It follows that $E\left[V_{j+1}^m(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j}) | \boldsymbol{\pi_j}\right]$ is concave in $\boldsymbol{\mu_j}$ as the positive combination of concave functions is concave. Therefore, we have $h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a\left(\hat{\mu} - \sum_{i=0}^{l} \mu_{j-1,k}\right)^+ - \gamma_a\left(\sum_{i=0}^{l} \mu_{j-1,k} - \hat{\mu}\right)^+ +$

$E\left[V^m_{j+1}(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j})|\boldsymbol{\pi_j}\right]$ is jointly concave in $(\boldsymbol{\mu_{j-1}}, \hat{\mu})$. For a jointly-concave function $f(\boldsymbol{x}, y)$ and a convex set $\mathcal{Y}$, we have $g(\boldsymbol{x}) = \max_{y\in\mathcal{Y}} f(\boldsymbol{x}, y)$ is concave in $\boldsymbol{x}$. Then it follows that

$$V^m_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) = \max_{\hat{\mu}\geq\sum_{i=1}^{l}\mu_{j-1,k}} \left\{ h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a\left(\hat{\mu} - \sum_{i=0}^{l}\mu_{j-1,k}\right)^+ - \gamma_a\left(\sum_{i=0}^{l}\mu_{j-1,k} - \hat{\mu}\right)^+ + E\left[V^m_{j+1}(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j})|\boldsymbol{\pi_j}\right]\right\}$$

is concave in $\boldsymbol{\mu_{j-1}}$.

We next show the optimal policy follows the control band structure. For $\sum_{i=1}^{l}\mu_{j-1,k} = 0$, we define

$$\underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) \triangleq \arg\max_{\hat{\mu}\in\mathbb{R}_+} \left\{h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) - c_a\hat{\mu} + E\left[V^m_{j+1}(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j})|\boldsymbol{\pi_j}\right]\right\} \tag{EC.91}$$

For any $\hat{\mu}_{j-1}$ that is sufficiently small, i.e., $\hat{\mu}_{j-1} < \underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$, it is optimal to adjust the capacity up to the level $\underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$.

For an arbitrary large $\hat{\mu}_{j-1}$, we define

$$\overline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) \triangleq \arg\max_{\hat{\mu}\geq\sum_{i=1}^{l}\mu_{j-1,k}} \left\{h_{j+l}(\boldsymbol{\pi_j}, \hat{\mu}) + \gamma_a\hat{\mu} + E\left[V^m_{j+1}(\boldsymbol{\Pi_{j+1}}, \boldsymbol{\mu_j})|\boldsymbol{\pi_j}\right]\right\} \tag{EC.92}$$

It is optimal for the firm to disinvest its capacity to $\overline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$ for all $\hat{\mu}_{j-1} > \overline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$. As $c_a \geq 0$ and $c_a \geq -\gamma_a$, it follows that $\overline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) \geq \underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$. Following the concavity of the value-to-go function, it is optimal for the firm to stay put when $\underline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}}) \leq \hat{\mu}_{j-1} \leq \overline{\mu}_j(\boldsymbol{\pi_j}, \boldsymbol{\mu_{j-1}})$. Therefore, we have proved the optimal policy is a control band policy. □

EC.3. Profit and cost parameter estimations in numerical examples.

We use *Production* to indicate the total production volume of Ford in 2012, which is approximated by the wholesale volume of 5,668 thousands units (operating highlights, Ford Motor Company 2012). As estimated by IHS Automotive (P.12, Ford Motor Company 2012), the global automotive industry production capacity for light vehicles is about 108 million units, which exceeds the global production by 26 million units. We therefore use the industry capacity utilization $Utilization = \frac{108-26}{108} = 75.93\%$ to estimated Ford's total capacity (including all types of products) in 2012 as

$$Capacity = \frac{Production}{Utilization} = \frac{5,668\times 10^3}{75.93\%} = 7,465\times 10^3 units/year = 622.1\times 10^3 units/month$$

- Capacity adjustment cost c_a and γ_a. The capacity adjustment cost c_a is estimated from the *Amortization of special tools*(AST) (P.102, Ford Motor Company 2012). As Ford generally amortizes special tools over the expected life of a product program using a straightline method, we calculate the expected cost to install one unit of capacity c_a as

$$c_a = \frac{AST\times\frac{T}{2}}{Capacity} = \frac{1,861\times 10^6\times\frac{3}{2}}{622.1\times 10^3} = 4,487\ dollars\cdot month/units.$$

As the capacity adjustment is often irreversible, we use a coefficient γ to measure the irreversibility and assume $\gamma_a = \gamma c_a$. In the base analysis, we assume $\gamma = 0.1$, i.e., it is costly for the firm to downsize its capacity.

- Capacity overhead cost c_0. The overhead cost is estimated from *Maintenance and rearrangement* expense (MR) (P.102, Ford Motor Company 2012). This cost reflects the firm's expense to conduct routine maintenance and repair to keep up its capacity level, and is incurred regardless of the production location. Therefore, we calculate c_0 as

$$c_0 = \frac{MR}{Capacity} = \frac{1,352 \times 10^6}{7,465 \times 10^3} = 181.1 \ dollars/units.$$

- Capacity outsourcing cost c_1. The outsourcing cost is incurred when the demand exceeds the installed capacity and therefore has to be satisfied by another facility. Therefore, the capacity outsourcing cost includes the cost to maintain the extra unit of outsourcing capacity and additional machine setup and transportation costs, and we denote the cost $c_1 = (1+\beta)c_0$ with $\beta > 0$. In the base case, we assume $\beta = 1$.

- Unit profit p. The unit profit is the profit the firm earns from selling a car, excluding the capacity related cost. We denote the gross revenue by *Revenue* and the total operating cost by $Cost$. Then we estimate the unit profit as

$$\begin{aligned} p &= \frac{Revenue - Cost + MR + AST}{Production} = \frac{125,567 - 121,584 + 1,352 + 1,861}{5,668} \times 10^3 \\ &= 1,270 \ dollars/units. \end{aligned}$$

We observe that from Ford Focus's official website[7], a simple average of the starting manufacturer suggested retail price (MSRP) for the seven current focus models yields a value of $(16,200 + 18,200 + 19,200 + 23,200 + 23,799 + 24,200 + 39,200)/7 = \$23,414$. We observe that this value is close to the average retail price estimated from the financial data, $Revenue/Production = 125,567/5,668 \times 10^3 = \$22,154$.

In the numerical analysis, we perform robustness checks with respect to these estimated parameters as follows.

Misspecified demand. The base case has assumed three demand types: low, medium and high. In the low demand scenario, we assume the average demand decreases to 12.21 thousand units per month. In the medium demand scenario, we assume that the demand remains at the same level as the demand for the first and second generation. In the high demand scenario, we assume the average demand has increased to 19.71 thousand units. However, these assumptions may not be accurate. Therefore, we now analyze the case when the firm has incorrect information about the demand type and resulting distribution. When calculating the deterministic upper bound, the firm still has complete information about the demand.

Table EC.1 Regret for the misspecified demand and cost parameters

Parameters	Regret			
	mean	stdev	min	max
low demand mean θ_l	6.02%	0.31%	5.78%	7.71%
high demand mean θ_h	5.81%	0.51%	4.32%	6.77%
outsourcing and overhead cost $\beta=(c_1-c_0)/c_0$	6.08%	0.60%	4.99%	7.71%
downsizing and expansion cost $\gamma=\gamma_a/c_a$	6.01%	0.30%	5.45%	6.92%

Note: The number of observations for each case is 147.

Note that our two-step heuristic does not depend on the firm's knowledge about the high or low type demand: The demand information is needed for evaluation and comparison only. In the analysis, we vary the average demand of low demand type from -20% to 40% in the increment of 10%, and similar for the high demand type. For each set of demand parameters, we also vary the prior as $(0.2i, 0.2j, 1-0.2i-0.2j)$ where $i=0,1,...,5$ and $j=0,1,...,5-i$. We summarize the test statistics in the first two rows of Table EC.1. We observe that the average regret with respect to the relaxed upper bound is only about 6% with a range less than 2.45%, which indicates the performance of the regret is quite robust with respect to the misspecified demand parameters.

Cost parameters. We also analyze the impact of the cost parameter changes on the two-step heuristic. In particular, we examine this by varying the relative difference between the outsourcing cost and capacity overhead cost: $\beta=(c_1-c_0)/c_0$ fixing c_0, and the ratio of the downsizing cost to the expansion cost: $\gamma=\gamma_a/c_a$ fixing c_a. In our base case, we have $\beta=(362.2-181.1)/181.1=1$ and $\gamma=448.7/4,487=0.1$ (see Table 3). Similarly to the misspecified demand scenario, we also vary the prior as $(0.2i, 0.2j, 1-0.2i-0.2j)$ where $i=0,1,...,5$ and $j=0,1,...,5-i$. In a quite broad range of β (from 0.7 to 1.3) and γ (from -0.3 to 0.3), the regret does not change in any significant manner (see the third and fourth row of Table EC.1). These results show that the heuristic is quite robust with respect to the cost parameters, as the increase in the regret is smaller than 2.72% when the cost parameters and the prior vary.

References

Chen, W., M. Dawande, G. Janakiraman. 2014. Fixed-dimensional stochastic dynamic programs: An approximation scheme and an inventory application. *Oper. Res.* **62**(1) 81–103.

Ford Motor Company. 2012. Profitable growth for all: Ford Motor Company 2012 annual report `http://corporate.ford.com/doc/ar2012-2012%20Annual%20Report.pdf`. Retrieved June 1, 2013.

Gallego, G. 1992. A minmax distribution free procedure for the (Q,R) inventory model. *Oper. Res. Lett.* **11**(1) 55–60.

Lovász, L. 1983. Submodular functions and convexity. *Mathematical Programming The State of the Art*. Springer, 235–257.

Murota, K. 2003. *Discrete convex analysis*. SIAM.

Semple, J. 2007. Note: Generalized notions of concavity with an application to capacity management. *Oper. Res.* **55**(2) 284–291.

Topkis, D. M. 1998. *Supermodularity and complementarity*. Princeton university press.

Ye, Q., I. Duenyas. 2007. Optimal capacity investment decisions with two-sided fixed-capacity adjustment costs. *Oper. Res.* **55**(2) 272–283.

Zipkin, P. 2008. On the structure of lost-sales inventory models. *Oper. Res.* **56**(4) 937–944.